



Anisotropic triangulations via discrete Riemannian Voronoi diagrams

Jean-Daniel Boissonnat, Mael Rouxel-Labbé, Mathijs Wintraecken

**RESEARCH
REPORT**

N° XXXX

March 2017

Project-Team Datashape



Anisotropic triangulations via discrete Riemannian Voronoi diagrams

Jean-Daniel Boissonnat, Mael Rouxel-Labbé, Mathijs Wintraecken

Project-Team Datashape

Research Report n° XXXX — March 2017 — 51 pages

Abstract: The construction of anisotropic triangulations is desirable for various applications, such as the numerical solving of partial differential equations and the representation of surfaces in graphics. To solve this notoriously difficult problem in a practical way, we introduce the discrete Riemannian Voronoi diagram, a discrete structure that approximates the Riemannian Voronoi diagram. This structure has been implemented and was shown to lead to good triangulations in \mathbb{R}^2 and on surfaces embedded in \mathbb{R}^3 as detailed in our experimental companion paper.

In this paper, we study theoretical aspects of our structure. Given a finite set of points \mathcal{P} in a domain Ω equipped with a Riemannian metric, we compare the discrete Riemannian Voronoi diagram of \mathcal{P} to its Riemannian Voronoi diagram. Both diagrams have dual structures called the discrete Riemannian Delaunay and the Riemannian Delaunay complex. We provide conditions that guarantee that these dual structures are identical. It then follows from previous results that the discrete Riemannian Delaunay complex can be embedded in Ω under sufficient conditions, leading to an anisotropic triangulation with curved simplices. Furthermore, we show that, under similar conditions, the simplices of this triangulation can be straightened.

Key-words: Riemannien geometry, Voronoi diagram, Delaunay triangulation

Long version of the paper accepted at the 33rd International Symposium on Computational Geometry (2017)

**RESEARCH CENTRE
SOPHIA ANTIPOLIS – MÉDITERRANÉE**

2004 route des Lucioles - BP 93
06902 Sophia Antipolis Cedex

Triangulations anisotropiques via diagrammes de Voronoi riemanniens discrets

Résumé : L'utilisation de triangulations anisotropes est souhaitable dans de nombreux domaines, tels que la résolution d'équations aux dérivées partielles ou la visualisation de surfaces. Pour résoudre ce problème notoirement difficile, nous proposons l'utilisation de diagrammes de Voronoi riemanniens discrets, une structure discrete qui approxime le diagramme de Voronoi riemannien. Cette structure a été implémentée et nous avons montré dans une publication empirique associée qu'elle produisait de bonnes triangulations pour des domaines de \mathbb{R}^2 et des surfaces plongées dans \mathbb{R}^3 .

Dans ce papier, nous étudions les aspects théoriques de notre structure. Etant donné un ensemble fini de points \mathcal{P} dans un domaine Ω équipé d'une métrique riemannienne, nous comparons le diagramme de Voronoi riemannien discret de \mathcal{P} à sa version exacte. Ces deux diagrammes ont chacun une structure duale, respectivement appelées le complexe de Delaunay riemannien discret et le complexe de Delaunay riemannien. Nous donnons des conditions qui garantissent que ces deux complexes sont identiques. Il en résulte de résultats précédemment établis que le complexe de Delaunay riemannien discret peut être plongé dans Ω sous certaines conditions et une triangulation anisotrope faite de simplexes courbes est obtenue. En outre, nous montrons que, sous des conditions analogues, les simplexes de cette triangulation peuvent être rendus droits.

Mots-clés : Géométrie riemannienne, Diagramme de Voronoi, Triangulation de Delaunay

1 Introduction

Anisotropic triangulations are triangulations whose elements are elongated along prescribed directions. Anisotropic triangulations are known to be well suited when solving PDE's [12, 22, 27]. They can also significantly enhance the accuracy of a surface representation if the anisotropy of the triangulation conforms to the curvature of the surface [18].

Many methods to generate anisotropic triangulations are based on the notion of Riemannian metric and create triangulations whose elements adapt locally to the size and anisotropy prescribed by the local geometry. The numerous theoretical and practical results [1] of the Euclidean Voronoi diagram and its dual structure, the Delaunay triangulation, have pushed authors to try and extend these well-established concepts to the anisotropic setting. Labelle and Shewchuk [20] and Du and Wang [14] independently introduced two anisotropic Voronoi diagrams whose anisotropic distances are based on a discrete approximation of the Riemannian metric field. Contrary to their Euclidean counterpart, the fact that the dual of these anisotropic Voronoi diagrams is an embedded triangulation is not immediate, and, despite their strong theoretical foundations, the anisotropic Voronoi diagrams of Labelle and Shewchuk and Du and Wang have only been proven to yield, under certain conditions, a good triangulation in a two-dimensional setting [8, 9, 11, 14, 20].

Both these anisotropic Voronoi diagrams can be considered as an approximation of the exact Riemannian Voronoi diagram, whose cells are defined as $V_g(p_i) = \{x \in \Omega \mid d_g(p_i, x) \leq d_g(p_j, x), \forall p_j \in \mathcal{P} \setminus \{p_i\}\}$, where $d_g(p, q)$ denotes the geodesic distance. Their main advantage is to ease the computation of the anisotropic diagrams. However, their theoretical and practical results are rather limited. The exact Riemannian Voronoi diagram comes with the benefit of providing a more favorable theoretical framework and recent works have provided sufficient conditions for a point set to be an embedded Riemannian Delaunay complex [2, 16, 21]. We approach the Riemannian Voronoi diagram and its dual Riemannian Delaunay complex with a focus on both practicality and theoretical robustness. We introduce the discrete Riemannian Voronoi diagram, a discrete approximation of the (exact) Riemannian Voronoi diagram. Experimental results, presented in our companion paper [26], have shown that this approach leads to good anisotropic triangulations for two-dimensional domains and surfaces, see Figure 1.

We introduce in this paper the theoretical side of this work, showing that our approach is theoretically sound in all dimensions. We prove that, under sufficient conditions, the discrete Riemannian Voronoi diagram has the same combinatorial structure as the (exact) Riemannian Voronoi diagram and that the dual discrete Riemannian Delaunay complex can be embedded as a triangulation of the point set, with either curved or straight simplices. Discrete Voronoi diagrams have been independently studied, although in a two-dimensional isotropic setting by Cao et al. [10].

2 Riemannian geometry

In the main part of the text we consider an (open) domain Ω in \mathbb{R}^n endowed with a Riemannian metric g , which we shall discuss below. We assume that the metric g is Lipschitz continuous. The structures of interest will be built from a finite set of points \mathcal{P} , which we call *sites*.

2.1 Riemannian metric

A *Riemannian metric field* g , defined over Ω , associates a *metric* $g(p) = G_p$ to any point p of the domain. This means that for any $v, w \in \mathbb{R}^n$ we associate an inner product $\langle v, w \rangle_g = v^t g(p) w$, in a way that smoothly depends on p . Using a Riemannian metric, we can associate lengths to




Figure 1: Left, the discrete Riemannian Voronoi diagram (colored cells with bisectors in white) and its dual complex (in black) realized with straight simplices of a two-dimensional domain endowed with a hyperbolic shock-based metric field. Right, the discrete Riemannian Voronoi diagram and the dual complex realized with curved simplices of the “chair” surface endowed with a curvature-based metric field [26].

curves and define the geodesic distance d_g as the minimizer of the lengths of all curves between two points. When the map $g : p \mapsto G$ is constant, the metric field is said to be *uniform*. In this case, the distance between two points x and y in Ω is $d_G(x, y) = \|x - y\|_G = \sqrt{(x - y)^t G (x - y)}$.

Most traditional geometrical objects can be generalized using the geodesic distance. For example, the geodesic (closed) ball centered on $p \in \Omega$ and of radius r is given by $B_g(p, r) = \{x \in \Omega \mid d_g(p, x) \leq r\}$. In the following, we assume that $\Omega \subset \mathbb{R}^n$ is endowed with a Lipschitz continuous metric field g .

We define the *metric distortion* between two distance functions $d_g(x, y)$ and $d_{g'}(x, y)$ to be the function $\psi(g, g')$ such that for all x, y in a small-enough neighborhood we have: $1/\psi(g, g') d_g(x, y) \leq d_{g'}(x, y) \leq \psi(g, g') d_g(x, y)$. Observe that $\psi(g, g') \geq 1$ and $\psi(g, g') = 1$ when $g = g'$. Our definition generalizes the concept of distortion between two metrics $g(p)$ and $g(q)$, as defined by Labelle and Shewchuk [20] (see Appendix B).

2.2 Geodesy

Let $v \in \mathbb{R}^n$. From the unique geodesic γ satisfying $\gamma(0) = p$ with initial tangent vector $\dot{\gamma} = v$, one defines the *exponential map* through $\exp(v) = \gamma(1)$. The *injectivity radius* at a point p of Ω is the largest radius for which the exponential map at p restricted to a ball of that radius is a diffeomorphism. The injectivity radius ι_Ω of Ω is defined as the infimum of the injectivity radii at all points. For any $p \in \Omega$ and for a two-dimensional linear subspace H of the tangent space at p , we define the *sectional curvature* K at p for H as the Gaussian curvature at p of the surface $\exp_p(H)$.

In the theoretical studies of our algorithm, we will assume that the injectivity radius of Ω is strictly positive and its sectional curvatures are bounded.

2.3 Power protected nets

Controlling the quality of the Delaunay and Voronoi structures will be essential in our proofs. For this purpose, we use the notions of net and of power protection.

Power protection of point sets Power protection of simplices is a concept formally introduced by Boissonnat, Dyer and Ghosh [2]. Let σ be a simplex whose vertices belong to \mathcal{P} , and let $B_g(\sigma) = B_g(c, r)$ denote a circumscribing ball of σ where $r = d_g(c, p)$ for any vertex p of σ . We call c the circumcenter of σ and r its circumradius.

For $0 \leq \delta \leq r$, we associate to $B_g(\sigma)$ the dilated ball $B_g^{+\delta}(\sigma) = B(c, \sqrt{r^2 + \delta^2})$. We say that σ is δ -power protected if $B_g^{+\delta}(\sigma)$ does not contain any point of $\mathcal{P} \setminus \text{Vert}(\sigma)$ where $\text{Vert}(\sigma)$ denotes the vertex set of σ . The ball $B_g^{+\delta}$ is the *power protected* ball of σ . Finally, a point set \mathcal{P} is δ -power protected if the Delaunay ball of its simplices are δ -power protected.

Nets To ensure that the simplices of the structures that we shall consider are well shaped, we will need to control the density and the sparsity of the point set. The concept of net conveys these requirements through *sampling* and *separation* parameters.

The sampling parameter is used to control the density of a point set: if Ω is a bounded domain, \mathcal{P} is said to be an ε -sample set for Ω with respect to a metric field g if $d_g(x, \mathcal{P}) < \varepsilon$, for all $x \in \Omega$. The sparsity of a point set is controlled by the separation parameter: the set \mathcal{P} is said to be μ -separated with respect to a metric field g if $d_g(p, q) \geq \mu$ for all $p, q \in \mathcal{P}$. If \mathcal{P} is an ε -sample that is μ -separated, we say that \mathcal{P} is an (ε, μ) -net.

3 Riemannian Delaunay triangulations

Given a metric field g , the *Riemannian Voronoi diagram* of a point set \mathcal{P} , denoted by $\text{Vor}_g(\mathcal{P})$, is the Voronoi diagram built using the geodesic distance d_g . Formally, it is a partition of the domain in *Riemannian Voronoi cells* $\{V_g(p_i)\}$, where $V_g(p_i) = \{x \in \Omega \mid d_g(p_i, x) \leq d_g(p_j, x), \forall p_j \in \mathcal{P} \setminus p_i\}$.

The Riemannian Delaunay complex of \mathcal{P} is an abstract simplicial complex, defined as the nerve of the Riemannian Voronoi diagram, that is the set of simplices $\text{Del}_g(\mathcal{P}) = \{\sigma \mid \text{Vert}(\sigma) \in \mathcal{P}, \cap_{p \in \sigma} V_g(p) \neq \emptyset\}$. There is a straightforward duality between the diagram and the complex, and between their respective elements.

In this paper, we will consider both abstract simplices and complexes, as well as their geometric realization in \mathbb{R}^n with vertex set \mathcal{P} . We now introduce two realizations of a simplex that will be useful, one curved and the other one straight.

The *straight realization* of a n -simplex σ with vertices in \mathcal{P} is the convex hull of its vertices. We denote it by $\bar{\sigma}$. In other words,

$$\bar{\sigma} = \{x \in \Omega \subset \mathbb{R}^n \mid x = \sum_{p \in \sigma} \lambda_p(x) p, \lambda_p(x) \geq 0, \sum_{p \in \sigma} \lambda_p(x) = 1\}. \quad (1)$$

The *curved realization*, noted $\tilde{\sigma}$ is based on the notion of Riemannian center of mass [19, 15]. Let y be a point of $\bar{\sigma}$ with barycentric coordinate $\lambda_p(y), p \in \sigma$. We can associate the energy functional $\mathcal{E}_y(x) = \frac{1}{2} \sum_{p \in \sigma} \lambda_p(y) d_g(x, p)^2$. We then define the curved realization of σ as

$$\tilde{\sigma} = \{\tilde{x} \in \Omega \subset \mathbb{R}^n \mid \tilde{x} = \text{argmin } \mathcal{E}_{\tilde{x}}(x), \tilde{x} \in \bar{\sigma}\}. \quad (2)$$

The edges of $\tilde{\sigma}$ are geodesic arcs between the vertices. Such a curved realization is well defined provided that the vertices of σ lie in a sufficiently small ball according to the following theorem of Karcher [19].

Theorem 3.1 (Karcher). *Let the sectional curvatures K of Ω be bounded, that is $\Lambda_- \leq K \leq \Lambda_+$. Let us consider the function \mathcal{E}_y on B_ρ , a geodesic ball of radius ρ that contains the set $\{p_i\}$.*

Assume that $\rho \in \mathbb{R}^+$ is less than half the injectivity radius and less than $\pi/4\sqrt{\Lambda_+}$ if $\Lambda_+ > 0$. Then \mathcal{E}_y has a unique minimum point in B_ρ , which is called the center of mass.

Given an (abstract) simplicial complex \mathcal{K} with vertices in \mathcal{P} , we define the straight (resp., curved) realization of \mathcal{K} as the collection of straight (resp., curved) realizations of its simplices, and we write $\bar{\mathcal{K}} = \{\bar{\sigma}, \sigma \in \mathcal{K}\}$ and $\tilde{\mathcal{K}} = \{\tilde{\sigma}, \sigma \in \mathcal{K}\}$.

We will consider the case where \mathcal{K} is $\text{Del}_g(\mathcal{P})$. A simplex of $\bar{\text{Del}}_g(\mathcal{P})$ will simply be called a straight Riemannian Delaunay simplex and a simplex of $\tilde{\text{Del}}_g(\mathcal{P})$ will be called a curved Riemannian Delaunay simplex, omitting “realization of”. In the next two sections, we give sufficient conditions for $\bar{\text{Del}}_g(\mathcal{P})$ and $\tilde{\text{Del}}_g(\mathcal{P})$ to be embedded in Ω , in which case we will call them the straight and the curved Riemannian *triangulations* of \mathcal{P} .

3.1 Sufficient conditions for $\tilde{\text{Del}}_g(\mathcal{P})$ to be a triangulation of \mathcal{P}

It is known that $\tilde{\text{Del}}_g(\mathcal{P})$ is embedded in Ω under sufficient conditions. We give a short overview of these results. As in Dyer et al. [15], we define the non-degeneracy of a simplex $\tilde{\sigma}$ of $\tilde{\text{Del}}_g(\mathcal{P})$.

Definition 3.2. *The curved realization $\tilde{\sigma}$ of a Riemannian Delaunay simplex σ is said to be non-degenerate if and only if it is homeomorphic to the standard simplex.*

Sufficient conditions for the complex $\tilde{\text{Del}}_g(\mathcal{P})$ to be embedded in Ω were given in [15]: a curved simplex is known to be non-degenerate if the Euclidean simplex obtained by lifting the vertices to the tangent space at one of the vertices via the exponential map has sufficient quality compared to the bounds on sectional curvature. Here, good quality means that the simplex is well shaped, which may be expressed either through its fatness (volume compared to longest edge length) or its thickness (smallest height compared to longest edge length).

Let us assume that, for each vertex p of $\text{Del}_g(\mathcal{P})$, all the curved Delaunay simplices in a neighborhood of p are non-degenerate and patch together well. Under these conditions, $\tilde{\text{Del}}_g(\mathcal{P})$ is embedded in Ω . We call $\tilde{\text{Del}}_g(\mathcal{P})$ the *curved Riemannian Delaunay triangulation* of \mathcal{P} .

3.2 Sufficient conditions for $\bar{\text{Del}}_g(\mathcal{P})$ to be a triangulation of \mathcal{P}

Assuming that the conditions for $\tilde{\text{Del}}_g(\mathcal{P})$ to be embedded in Ω are satisfied, we now give conditions such that $\bar{\text{Del}}_g(\mathcal{P})$ is also embedded in Ω . The key ingredient will be a bound on the distance between a point of a simplex $\tilde{\sigma}$ and the corresponding point on the associated straight simplex $\bar{\sigma}$ (Lemma 3.3). This bound depends on the properties of the set of sites and on the local distortion of the metric field. When this bound is sufficiently small, $\bar{\text{Del}}_g(\mathcal{P})$ is embedded in Ω as stated in Theorem 3.4.

Lemma 3.3. *Let σ be an n -simplex of $\text{Del}_g(\mathcal{P})$. Let \bar{x} be a point of $\bar{\sigma}$ and \tilde{x} the associated point on $\tilde{\sigma}$ (as defined in Equation 1). If the geodesic distance d_g is close to the Euclidean distance $d_{\mathbb{E}}$, i.e. the distortion $\psi(g, g_{\mathbb{E}})$ is bounded by ψ_0 , then $|\tilde{x} - \bar{x}| \leq \sqrt{2 \cdot 4^3(\psi_0 - 1)\varepsilon^2}$.*

We now apply Lemma 3.3 to the facets of the simplices of $\tilde{\text{Del}}_g(\mathcal{P})$. The altitude of the vertex p in a simplex τ is noted $D(p, \tau)$.

Theorem 3.4. *Let \mathcal{P} be a δ -power protected (ε, μ) -net with respect to g on Ω . Let σ be any n -simplex of $\text{Del}_g(\mathcal{P})$ and p be any vertex of σ . Let τ be a facet of σ opposite of vertex p . If, for all $\tilde{x} \in \tilde{\tau}$, we have $|\tilde{x} - \bar{x}| \leq D(p_i, \sigma)$ (\bar{x} is defined in Equation 1), then $\bar{\text{Del}}_g(\mathcal{P})$ is embedded in Ω .*

The condition $|\tilde{x} - \bar{x}| \leq D(p_i, \sigma)$ is achieved for a sufficiently dense sampling according to Lemma 3.3 and the fact that the distortion $\psi_0 = \psi(g, g_{\mathbb{E}})$ goes to 1 when the density increases. The complete proofs of Lemma 3.3 and Theorem 3.4 can be found in Appendix F.

4 Discrete Riemannian structures

Although Riemannian Voronoi diagrams and Delaunay triangulations are appealing from a theoretical point of view, they are very difficult to compute in practice despite many studies [24]. To circumvent this difficulty, we introduce the discrete Riemannian Voronoi diagram. This discrete structure is easy to compute (see our companion paper [26] for details) and, as will be shown in the following sections, it is a good approximation of the exact Riemannian Voronoi diagram. In particular, their dual Delaunay structures are identical under appropriate conditions.

We assume that we are given a dense triangulation of the domain Ω we call the *canvas* and denote by \mathcal{C} . The canvas will be used to approximate geodesic distances between points of Ω and to construct the discrete Riemannian Voronoi diagram of \mathcal{P} , which we denote by $\text{Vor}_g^d(\mathcal{P})$. This bears some resemblance to the graph-induced complex of Dey et al. [13]. Notions related to the canvas will explicitly carry *canvas* in the name (for example, an edge of \mathcal{C} is a *canvas edge*). In our analysis, we shall assume that the canvas is a dense triangulation, although weaker and more efficient structures can be used (see Section 9 and [26]).

4.1 The discrete Riemannian Voronoi Diagram

To define the discrete Riemannian Voronoi diagram of \mathcal{P} , we need to give a unique color to each site of \mathcal{P} and to color the vertices of the canvas accordingly. Specifically, each canvas vertex is colored with the color of its closest site.

Definition 4.1 (Discrete Riemannian Voronoi diagram). *Given a metric field g , we associate to each site p_i its discrete cell $V_g^d(p_i)$ defined as the union of all canvas simplices with at least one vertex of the color of p_i . We call the set of these cells the discrete Riemannian Voronoi diagram of \mathcal{P} , and denote it by $\text{Vor}_g^d(\mathcal{P})$.*

Observe that contrary to typical Voronoi diagrams, our discrete Riemannian Voronoi diagram is not a partition of the canvas. Indeed, there is a one canvas simplex-thick overlapping since each canvas simplex $\sigma_{\mathcal{C}}$ belongs to all the Voronoi cells whose sites' colors appear in the vertices of $\sigma_{\mathcal{C}}$. This is intentional and allows for a straightforward definition of the complex induced by this diagram, as shown below.

4.2 The discrete Riemannian Delaunay complex

We define the *discrete Riemannian Delaunay complex* as the set of simplices $\text{Del}_g^d(\mathcal{P}) = \{\sigma \mid \text{Vert}(\sigma) \in \mathcal{P}, \cap_{p \in \sigma} V_g^d(p) \neq \emptyset\}$. Using a triangulation as canvas offers a very intuitive way to construct the discrete complex since each canvas k -simplex σ of \mathcal{C} has $k+1$ vertices $\{v_0, \dots, v_k\}$ with respective colors $\{c_0, \dots, c_k\}$ corresponding to the sites $\{p_{c_0}, \dots, p_{c_k}\} \in \mathcal{P}$. Due to the way discrete Voronoi cells overlap, a canvas simplex $\sigma_{\mathcal{C}}$ belongs to each discrete Voronoi cell whose color appears in the vertices of σ . Therefore, the intersection of the discrete Voronoi cells $\{V_g^d(p_i)\}$ whose colors appear in the vertices of σ is non-empty and the simplex σ with vertices $\{p_i\}$ thus belongs to the discrete Riemannian Delaunay complex. In that case, we say that the canvas simplex $\sigma_{\mathcal{C}}$ *witnesses* (or is a witness of) σ . For example, if the vertices of a canvas 3-simplex $\tau_{\mathcal{C}}$ have colors yellow–blue–blue–yellow, then the intersection of the discrete Voronoi cells

of the sites p_{yellow} and p_{blue} is non-empty and the one-simplex σ with vertices p_{yellow} and p_{blue} belongs to the discrete Riemannian Delaunay complex. The canvas simplex τ_C thus witnesses the (abstract, for now) edge between p_{yellow} and p_{blue} .

Figure 2 illustrates a canvas painted with discrete Voronoi cells, and the witnesses of the discrete Riemannian Delaunay complex.




Figure 2: A canvas (black edges) and a discrete Riemannian Voronoi diagram drawn on it. The canvas simplices colored in red are witnesses of Voronoi vertices. The canvas simplices colored in grey are witnesses of Voronoi edges. Canvas simplices whose vertices all have the same color are colored with that color.

Remark 4.2. *If the intersection $\bigcap_{i=0\dots k} V_g^d(p_{c_i})$ is non-empty, then the intersection of any subset of $\{V_g^d(p_{c_i})\}_{i=0\dots k}$ is non-empty. In other words, if a canvas simplex σ_C witnesses a simplex σ , then for each face τ of σ , there exists a face τ_C of σ_C that witnesses τ . As we assume that there is no boundary, the complex is pure and it is sufficient to only consider canvas n -simplices whose vertices have all different colors to build $\text{Del}_g^d(\mathcal{P})$.*

Similarly to the definition of curved and straight Riemannian Delaunay complexes, we can define their discrete counterparts we respectively denote by $\widetilde{\text{Del}}_g^d(\mathcal{P})$ and $\overline{\text{Del}}_g^d(\mathcal{P})$. We will now exhibit conditions such that these complexes are well-defined and embedded in Ω .

5 Equivalence between the discrete and the exact structures

We first give conditions such that $\text{Vor}_g^d(\mathcal{P})$ and $\text{Vor}_g(\mathcal{P})$ have the same combinatorial structure, or, equivalently, that the dual Delaunay complexes $\text{Del}_g(\mathcal{P})$ and $\text{Del}_g^d(\mathcal{P})$ are identical. Under these conditions, the fact that $\text{Del}_g^d(\mathcal{P})$ is embedded in Ω will immediately follow from the fact that the exact Riemannian Delaunay complex $\text{Del}_g(\mathcal{P})$ is embedded (see Sections 3.1 and 3.2). It thus remains to exhibit conditions under which $\text{Del}_g^d(\mathcal{P})$ and $\text{Del}_g(\mathcal{P})$ are identical.

Requirements will be needed on both the set of sites in terms of density, sparsity and protection, and on the density of the canvas. The central idea in our analysis is that power protection of \mathcal{P} will imply a lower bound on the distance separating two non-adjacent Voronoi objects (and in particular two Voronoi vertices). From this lower bound, we will obtain an upper bound on the size on the cells of the canvas so that the combinatorial structure of the discrete diagram is the same as that of the exact one. The density of the canvas is expressed by e_C , the length of its longest edge.

The main result of this paper is the following theorem.

Theorem 5.1. *Assume that \mathcal{P} is a δ -power protected (ε, μ) -net in Ω with respect to g . Assume further that ε is sufficiently small and δ is sufficiently large compared to the distortion between $g(p)$ and g in an ε -neighborhood of p . Let $\{\lambda_i\}$ be the eigenvalues of $g(p)$ and ℓ_0 a value that depends on ε and δ (Precise bounds for ε, δ and ℓ_0 are given in the proof). Then, if $e_c < \min_{p \in \mathcal{P}} \left[\min_i (\sqrt{\lambda_i}) \min \{\mu/3, \ell_0/2\} \right]$, $\text{Del}_g^d(\mathcal{P}) = \text{Del}_g(\mathcal{P})$.*

The rest of the paper will be devoted to the proof of this theorem. Our analysis is divided into two parts. We first consider in Section 6 the most basic case of a domain of \mathbb{R}^n endowed with the Euclidean metric field. The result is given by Theorem 6.1. The assumptions are then relaxed and we consider the case of an arbitrary metric field over Ω in Section 7. As we shall see, the Euclidean case already contains most of the difficulties that arise during the proof and the extension to more complex settings will be deduced from the Euclidean case by bounding the distortion.

6 Equality of the Riemannian Delaunay complexes in the Euclidean setting

In this section, we restrict ourselves to the case where the metric field is the Euclidean metric $g_{\mathbb{E}}$. To simplify matters, we initially assume that geodesic distances are computed exactly on the canvas. The following theorem gives sufficient conditions to have equality of the complexes.

Theorem 6.1. *Assume that \mathcal{P} is a δ -power protected (ε, μ) -net of Ω with respect to the Euclidean metric field $g_{\mathbb{E}}$. Denote by \mathcal{C} the canvas, a triangulation with maximal edge length e_c . If $e_c < \min \{\mu/16, \delta^2/64\varepsilon\}$, then $\text{Del}_{\mathbb{E}}^d(\mathcal{P}) = \text{Del}_{\mathbb{E}}(\mathcal{P})$.*

We shall now prove Theorem 6.1 by enforcing the two following conditions which, combined, give the equality between the discrete Riemannian Delaunay complex and the Riemannian Delaunay complex:

- (1) for every Voronoi vertex in the Riemannian Voronoi diagram $v = \cap_{\{p_i\}} V_g(p_i)$, there exists at least one canvas simplex with the corresponding colors $\{c_{p_i}\}$;
- (2) no canvas simplex witnesses a simplex that does not belong to the Riemannian Delaunay complex (equivalently, no canvas simplex has vertices whose colors are those of non-adjacent Riemannian Voronoi cells).

Condition (2) is a consequence of the separation of Voronoi objects, which in turn follows from power protection. The separation of Voronoi objects has previously been studied, for example by Boissonnat et al. [2]. Although the philosophy is the same, our setting is slightly more difficult and the results using power protection are new and use a more geometrical approach (see Appendix C).

6.1 Sperner's lemma


Rephrasing Condition (1), we seek requirements on the density of the canvas \mathcal{C} and on the nature of the point set \mathcal{P} such that there exists at least one canvas n -simplex of \mathcal{C} that has exactly the colors c_0, \dots, c_d of the vertices p_0, \dots, p_d of a simplex σ , for all $\sigma \in \text{Del}_g(\mathcal{P})$. To prove the existence of such a canvas simplex, we employ Sperner's lemma [28], which is a discrete analog of Brouwer's fixed point theorem. We recall this result in Theorem 6.2 and illustrate it in a two-dimensional setting (inset).

Theorem 6.2 (Sperner's lemma).

Let $\sigma = (p_0, \dots, p_n)$ be an n -simplex and let T_σ denote a triangulation of the simplex. Let each vertex $v' \in T_\sigma$ be colored such that the following conditions are satisfied:

- The vertices p_i of σ all have different colors.
- If a vertex p' lies on a k -face $(p_{i_0}, \dots, p_{i_k})$ of σ , then p' has the same color as one of the vertices of the face, that is p_{i_j} .

Then, there exists an odd number of simplices in T_σ whose vertices are colored with all $n + 1$ colors. In particular, there must be at least one.



We shall apply Sperner's lemma to the canvas \mathcal{C} and show that for every Voronoi vertex v in the Riemannian Voronoi diagram, we can find a subset \mathcal{C}_v of the canvas that fulfills the assumptions of Sperner's lemma, hence obtaining the existence of a canvas simplex in \mathcal{C}_v (and therefore in \mathcal{C}) that witnesses σ_v . Concretely, the subset \mathcal{C}_v is obtained in two steps:

- We first apply a barycentric subdivision of the Riemannian Voronoi cells incident to v . From the resulting set of simplices, we extract a triangulation \mathcal{T}_v composed of the simplices incident to v (Section 6.2).
- We then construct the subset \mathcal{C}_v by overlaying the border of \mathcal{T}_v and the canvas (Section 6.3).

We then show that if the canvas simplices are small enough – in terms of edge length – then \mathcal{C}_v is the triangulation of a simplex that satisfies the assumptions of Sperner's lemma.

The construction of \mathcal{C}_v is detailed in the following sections and illustrated in Figure 3: starting from a colored canvas (left), we subdivide the incident Voronoi cells of v to obtain \mathcal{T}_v (middle), and deduce the set of canvas simplices \mathcal{C}_v which forms a triangulation that satisfies the hypotheses of Sperner's lemma, thus giving the existence of a canvas simplex (in green, right) that witnesses the Voronoi vertex within the union of the simplices, and therefore in the canvas.




Figure 3: Illustration of the construction of \mathcal{C}_v . The Riemannian Voronoi diagram is drawn with thick orange edges and the sites are colored squares. The canvas is drawn with thin gray edges and colored circular vertices. The middle frame shows the subdivision of the incident Voronoi cells with thick black edges and the triangulation \mathcal{T}_v is drawn in yellow. On the right frame, the set of simplices \mathcal{C}_v is colored in purple (simplices that do not belong to \mathcal{C}) and in dark yellow (simplices that belong to \mathcal{C}).

6.2 The triangulation \mathcal{T}_v

For a given Voronoi vertex v in the Euclidean Voronoi diagram $\text{Vor}_{\mathbb{E}}(\mathcal{P})$ of the domain Ω , the initial triangulation \mathcal{T}_v is obtained by applying a combinatorial barycentric subdivision of the Voronoi cells of $\text{Vor}_{\mathbb{E}}(\mathcal{P})$ that are incident to v : to each Voronoi cell V incident to v , we associate to each face F of V a point c_F in F which is not necessarily the geometric barycenter. We randomly associate to c_F the color of any of the sites whose Voronoi cells intersect to give F . For example, in a two-dimensional setting, if the face F is a Voronoi edge that is the intersection of V_{red} and V_{blue} , then c_F is colored either red or blue. Then, the subdivision of V is computed by associating to all possible sequences of faces $\{F_0, F_1, \dots, F_{n-1}, F_n\}$ such that $F_0 \subset F_1 \subset \dots \subset F_n = V$ and $\dim(F_{i+1}) = \dim(F_i) + 1$ the simplex with vertices $\{c_{F_0}, c_{F_1}, \dots, c_{F_{n-1}}, c_{F_n}\}$. These barycentric subdivisions are allowed since Voronoi cells are convex polytopes.

Denote by Σ_V the set of simplices obtained by barycentric subdivision of V and $\Sigma_v = \{\cup \Sigma_V \mid v \in V\}$. The triangulation \mathcal{T}_v is defined as the star of v in Σ_v , that is the set of simplices in Σ_v that are incident to v . \mathcal{T}_v is illustrated in Figure 4 in dimension 3. As shall be proven in Lemma 6.3, \mathcal{T}_v can be used to define a combinatorial simplex that satisfies the assumptions of Sperner's lemma.

\mathcal{T}_v as a triangulation of an n -simplex

By construction, the triangulation \mathcal{T}_v is a triangulation of the (Euclidean) Delaunay simplex σ_v dual of v as follows. We first perform the standard barycentric subdivision on this Delaunay simplex σ_v . We then map the barycenter of a k -face τ of σ_v to the point c_{F_i} on the Voronoi face F_i , where F_i is the Voronoi dual of the k -face τ . This gives a piecewise linear homeomorphism from the Delaunay simplex σ_v to the triangulation \mathcal{T}_v . We call the image of this map the simplex σ_S and refer to the images of the faces of the Delaunay simplex as the faces of σ_S . We can now apply Sperner's lemma.




Figure 4: The triangulation \mathcal{T}_v in 3D. A face (in

Lemma 6.3. *Let \mathcal{P} be a δ -power protected (ε, μ) -net. Let v be a Voronoi vertex in the Euclidean Voronoi diagram, $\text{Vor}_{\mathbb{E}}(\mathcal{P})$, and let Σ_v be defined as above. The simplex σ_S and the triangulation \mathcal{T}_v satisfy the assumptions of Sperner's lemma in dimension n .*

Proof. By the piecewise linear map that we have described above, \mathcal{T}_v is a triangulation of the simplex σ_S . Because by construction the vertices c_{F_i} lie on the Voronoi duals F_i of the corresponding Delaunay face τ , c_{F_i} has the one of the colors of the Delaunay vertices of τ . Therefore, σ_S satisfies the assumptions of Sperner's lemma and there exists an n -simplex in \mathcal{T}_v that witnesses v and its corresponding simplex σ_v in $\text{Del}_g(\mathcal{P})$. \square

6.3 Building the triangulation \mathcal{C}_v

Let p_i be the vertices of the k -face τ_S of σ_S . In this section we shall assume not only that τ_S is contained in the union of the Voronoi cells of $V(p_i)$, but in fact that τ_S is a distance $8e_C$ removed from the boundary of $\cup V(p_i)$, where e_C is the longest edge length of a simplex in the canvas. We will now construct a triangulation \mathcal{C}_v of σ_S such that:

- σ_S and its triangulation \mathcal{C}_v satisfy the conditions of Sperner's lemma,
- the simplices of \mathcal{C}_v that have no vertex that lies on the boundary $\partial\sigma_S$ are simplices of the canvas \mathcal{C} .

The construction goes as follows. We first intersect the canvas \mathcal{C} with σ_S and consider the canvas simplices $\sigma_{\mathcal{C},i}$ such that the intersection of σ_S and $\sigma_{\mathcal{C},i}$ is non-empty. These simplices $\sigma_{\mathcal{C},i}$ can be subdivided into two sets, namely those that lie entirely in the interior of σ_S , which we denote by $\sigma_{\mathcal{C},i}^{\text{int}}$, and those that intersect the boundary, denoted by $\sigma_{\mathcal{C},i}^{\partial}$.

The simplices $\sigma_{\mathcal{C},i}^{\text{int}}$ are added to the set \mathcal{C}_v . We intersect the simplices $\sigma_{\mathcal{C},i}^{\partial}$ with σ_S and triangulate the intersection. Note that $\sigma_{\mathcal{C},i}^{\partial} \cap \sigma_S$ is a convex polyhedron and thus triangulating it is not a difficult task. The vertices of the simplices in the triangulation of $\sigma_{\mathcal{C},i}^{\partial} \cap \sigma_S$ are colored according to which Voronoi cell they belong to. Finally, the simplices in the triangulation of $\sigma_{\mathcal{C},i}^{\partial} \cap \sigma_S$ are added to the set \mathcal{C}_v .

Since \mathcal{T}_v is a triangulation of σ_S , the set \mathcal{C}_v is by construction also a triangulation of σ_S . This triangulation trivially gives a triangulation of the faces τ_S . Because we assume that τ_S is contained in the union of its Voronoi cells, with a margin of $8\epsilon_C$ we now can draw two important conclusions:

- The vertices of the triangulation of each face τ_S have the colors of the vertices p_i of τ_S .
- None of the simplices in the triangulation of $\sigma_{\mathcal{C},i}^{\partial} \cap \sigma_S$ can have $n+1$ colors, because every such simplex must be close to one face τ_S , which means that it must be contained in the union of the Voronoi cells $V(p_i)$ of the vertices of τ_S .

We can now invoke Sperner's lemma; \mathcal{C}_v is a triangulation of the simplex σ_S whose every face has been colored with the appropriate colors (since σ_S triangulated by \mathcal{T}_v satisfies the assumptions of Sperner's lemma, see Lemma 6.3). This means that there is a simplex \mathcal{C}_v that is colored with $n+1$ colors. Because of our second observation above, the simplex with these $n+1$ colors must lie in the interior of σ_S and is thus a canvas simplex.

We summarize by the following lemma:

Lemma 6.4. *If every face τ_S of σ_S with vertices p_i is at distance $8\epsilon_C$ from the boundary of the union of its Voronoi cells $\partial(\cup V(p_i))$, then there exists a canvas simplex in \mathcal{C}_v such that it is colored with the same vertices as the vertices of σ_S .*

The key task that we now face is to guarantee that faces τ_S indeed lie well inside of the union of the appropriate Voronoi regions. This requires first and foremost power protection. Indeed, if a point set is power protected, the distance between a Voronoi vertex c and the Voronoi faces that are not incident to c , which we will refer to from now on as *foreign* Voronoi faces, can be bounded, as shown in the following Lemma:

Lemma 6.5. *Suppose that c is the circumcenter of a δ -power protected simplex σ of a Delaunay triangulation built from an ϵ -sample, then all foreign Voronoi faces are at least $\delta^2/8\epsilon$ far from c .*

The proof of this Lemma is given in the appendix (Section C.2).

In almost all cases, this result gives us the distance bound we require: we can assume that vertices $\{c_{F_0}, c_{F_1}, \dots, c_{F_{n-1}}, c_{F_n}\}$ which we used to construct \mathcal{T}_v , are well placed, meaning that there is a minimum distance between these vertices and foreign Voronoi objects. However it can still occur that foreign Voronoi objects are close to a face τ_S of σ_S . This occurs even in two dimensions, where a Voronoi vertex v' can be very close to a face τ_S because of obtuse angles, as illustrated in Figure 5.




Figure 5: The point v' can be arbitrarily close to \mathcal{T}_v , as shown by the red segments (left and center). After piecewise linear deformation, this issue is resolved, as seen by the green segments (right).

Thanks to power protection, we know that v' is removed from foreign Voronoi objects. This means that we can deform $\sigma_{\mathcal{S}}$ (in a piecewise linear manner) in a neighborhood of v' such that the distance between v' and all the faces of the deformed $\sigma_{\mathcal{S}}$ is lower bounded.

In general the deformation of $\sigma_{\mathcal{S}}$ is performed by “radially pushing” simplices away from the foreign Voronoi faces of v with a ball of radius $r = \min \{\mu/16, \delta^2/64\epsilon\}$. The value $\mu/16$ is chosen so that we do not move any vertex of σ_v (the dual of v): indeed, \mathcal{P} is μ -separated and thus $d_{\mathbb{E}}(p_i, p_j) > \mu$. The value $\delta^2/64\epsilon$ is chosen so that $\sigma_{\mathcal{S}}$ and its deformation stay isotopic (no “pinching” can happen), using Lemma 6.5. In fact it is advisable to use a piecewise linear version of “radial pushing”, to ensure that the deformation of $\sigma_{\mathcal{S}}$ is a polyhedron. This guarantees that we can triangulate the intersection, see Chapter 2 of Rourke and Sanderson [25]. After this deformation we can follow the steps we have given above to arrive at a well-colored simplex.

Lemma 6.6. *Let \mathcal{P} be a δ -power protected (ϵ, μ) -net. Let v be a Voronoi vertex of the Euclidean Voronoi diagram $\text{Vor}_{\mathbb{E}}(\mathcal{P})$, and \mathcal{T}_v as defined above. If the length e_c of the longest canvas edge is bounded as follows: $e_c < \min \{\mu/16, \delta^2/64\epsilon\}$, then there exists a canvas simplex that witnesses v and the corresponding simplex σ_v in $\text{Del}_{\mathbb{E}}(\mathcal{P})$.*

Conclusion So far, we have only proven that $\text{Del}_g(\mathcal{P}) \subseteq \text{Del}_g^d(\mathcal{P})$. The other inclusion, which corresponds to Condition (2) mentioned above, is much simpler: as long as a canvas edge is shorter than the smallest distance between a Voronoi vertex and a foreign face of the Riemannian Voronoi diagram, then no canvas simplex can witness a simplex that is not in $\text{Del}_g(\mathcal{P})$. Such a bound is already given by Lemma 6.5 and thus, if $e_c < \delta^2/8\epsilon$ then $\text{Del}_g^d(\mathcal{P}) \subseteq \text{Del}_g(\mathcal{P})$. Observe that this requirement is weaker than the condition imposed in Lemma 6.6 and it was thus already satisfied. It follows that $\text{Del}_g^d(\mathcal{P}) = \text{Del}_g(\mathcal{P})$ if $e_c < \min \{\mu/16, \delta^2/64\epsilon\}$, which concludes the proof of Theorem 6.1.

Remark 6.7. *Assuming that the point set is a δ -power protected (ϵ, μ) -net might seem like a strong assumption. However, it should be observed that any non-degenerate point set can be seen as a δ -power protected (ϵ, μ) -net, for a sufficiently large value of ϵ and sufficiently small values of δ and μ . Our results are therefore always applicable but the necessary canvas density increases as the quality of the point set worsens (Lemma 6.6). In our companion practical paper [26, Section 7], we showed how to generate δ -power protected (ϵ, μ) -nets for given values of ϵ , μ and δ .*

7 Extension to more complex settings

In the previous section, we have placed ourselves in the setting of an (open) domain endowed with the Euclidean metric field. To prove Theorem 5.1, we need to generalize Theorem 6.1 to more general metrics, which will be done in the two following subsections.

The common path to prove $\text{Del}_g^d(\mathcal{P}) = \text{Del}_g(\mathcal{P})$ in all settings is to assume that \mathcal{P} is a power protected net with respect to the metric field. We then use the stability of entities under small metric perturbations to take us back to the now solved case of the domain Ω endowed with an Euclidean metric field. Separation and stability of Delaunay and Voronoi objects has previously been studied by Boissonnat et al. [2, 5], but our work lives in a slightly more complicated setting. Moreover, our proofs are generally more geometrical and sometimes simpler. For completeness, the extensions of these results to our context are detailed in Appendix C for separation, and in Appendix E for stability.

We now detail the different intermediary settings. For completeness, the full proofs are included in the appendices.

7.1 Uniform metric field

We first consider the rather easy case of a non-Euclidean but uniform (constant) metric field over an (open) domain. The square root of a metric gives a linear transformation between the base space where distances are considered in the metric and a *metric space* where the Euclidean distance is used (see Appendix B.1.1). Additionally, we show that a δ -power protected (ε, μ) -net with respect to the uniform metric is, after transformation, still a δ -power protected (ε, μ) -net but with respect to the Euclidean setting (Lemma E.1 in Appendix E), bringing us back to the setting we have solved in Section 6. Bounds on the power protection, sampling and separation coefficients, and on the canvas edge length can then be obtained from the result for the Euclidean setting, using Theorem 6.6. These bounds can be transported back to the case of uniform metric fields by scaling these values according to the smallest eigenvalue of the metric (Theorem H.1 in Appendix H).

7.2 Arbitrary metric field

The case of an arbitrary metric field over Ω is handled by observing that an arbitrary metric field is locally well-approximated by a uniform metric field. It is then a matter of controlling the distortion.

We first show that, for any point $p \in \Omega$, density separation and power protection are locally preserved in a neighborhood U_p around p when the metric field g is approximated by the constant metric field $g' = g(p)$ (Lemmas E.2 and E.16 in Appendix E): if \mathcal{P} is a δ -power protected (ε, μ) -net with respect to g , then \mathcal{P} is a δ' -power protected (ε', μ') -net with respect to g' . Previous results can now be applied to obtain conditions on δ' , ε' , μ' and on the (local) maximal length of the canvas such that $\text{Del}_g^d(\mathcal{P}) = \text{Del}_g(\mathcal{P})$ (Lemma H.2 in Appendix H).

These local triangulations can then be stitched together to form a triangulation embedded in Ω . The (global) bound on the maximal canvas edge length is given by the minimum of the local bounds, each computed through the results of the previous sections. This ends the proof of Theorem 5.1.

Once the equality between the complexes is obtained, conditions giving the embeddability of the discrete Karcher Delaunay triangulation and the discrete straight Delaunay triangulation are given by previous results that we have established in Sections 3.1 and 3.2 respectively.

8 Extensions of the main result

Approximate geodesic computations Approximate geodesic distance computations can be incorporated in the analysis of the previous section by observing that computing inaccurately geodesic distances in a domain Ω endowed with a metric field g can be seen as computing exactly geodesic distances in Ω with respect to a metric field g' that is close to g (Section H.3 in Appendix H).

General manifolds The previous section may also be generalized to an arbitrary smooth n -manifold \mathcal{M} embedded in \mathbb{R}^m . We shall assume that, apart from the metric induced by the embedding of the domain in Euclidean space, there is a second metric g defined on \mathcal{M} . Let $\pi_p : \mathcal{M} \rightarrow T_p\mathcal{M}$ be the orthogonal projection of points of \mathcal{M} on the tangent space $T_p\mathcal{M}$ at p . For a sufficiently small neighborhood $U_p \subset T_p\mathcal{M}$, π_p is a local diffeomorphism (see Niyogi [23]).

Denote by \mathcal{P}_{T_p} the point set $\{\pi_p(p_i), p_i \in \mathcal{P}\}$ and \mathcal{P}_{U_p} the restriction of \mathcal{P}_{T_p} to U_p . Assuming that the conditions of Niyogi et al. [23] are satisfied (which are simple density constraints on ε compared to the reach of the manifold), the pullback of the metric with the inverse projection $(\pi_p^{-1})^*g$ defines a metric g_p on U_p such that for all $q, r \in U_p$, $d_{g_p}(q, r) = d_g(\pi_p^{-1}(q), \pi_p^{-1}(r))$. This implies immediately that if \mathcal{P} is a δ -power protected (ε, μ) -net on \mathcal{M} with respect to g then \mathcal{P}_{U_p} is a δ -power protected (ε, μ) -net on U_p . We have thus a metric on a subset of a n -dimensional space, in this case the tangent space, giving us a setting that we have already solved. It is left to translate the sizing field requirement from the tangent plane to the manifold \mathcal{M} itself. Note that the transformation π_p is completely independent of g . Boissonnat et al. [2, Lemma 3.7] give bounds on the metric distortion of the projection on the tangent space. This result allows to carry the canvas sizing field requirement from the tangent space to \mathcal{M} .

9 Implementation

The construction of the discrete Riemannian Voronoi diagram and of the discrete Riemannian Delaunay complex has been implemented for $n = 2, 3$ and for surfaces of \mathbb{R}^3 . An in-depth description of our structure and its construction as well as an empirical study can be found in our practical paper [26]. We simply make a few observations here.

The theoretical bounds on the canvas edge length provided by Theorems 5.1 and 6.1 are far from tight and thankfully do not need to be honored in practice. A canvas whose edge length are about a tenth of the distance between two seeds suffices. This creates nevertheless unnecessarily dense canvasses since the density does not in fact need to be equal everywhere at all points and even in all directions. This issue is resolved by the use of anisotropic canvasses.

Our analysis was based on the assumption that all canvas vertices are painted with the color of the closest site. In our implementation, we color the canvas using a multiple-front vector Dijkstra algorithm [7], which empirically does not suffer from the same convergence issues as the traditional Dijkstra algorithm, starting from all the sites. It should be noted that any geodesic distance computation method can be used, as long as it converges to the exact geodesic distance when the canvas becomes denser. The Riemannian Delaunay complex is built on the fly during the construction of the discrete Riemannian Voronoi diagram: when a canvas simplex is first fully colored, its combinatorial information is extracted and the corresponding simplex is added to $\text{Del}_g(\mathcal{P})$.

Acknowledgments

We thank Ramsay Dyer for enlightening discussions. The first and third authors have received funding from the European Research Council under the European Union’s ERC Grant Agreement number 339025 GUDHI (Algorithmic Foundations of Geometric Understanding in Higher Dimensions).

References

- [1] AURENHAMMER, F., AND KLEIN, R. Voronoi diagrams. In *Handbook of Computational Geometry*, J. Sack and G. Urrutia, Eds. Elsevier Science Publishing, 2000, pp. 201–290.
- [2] BOISSONNAT, J.-D., DYER, R., AND GHOSH, A. Delaunay triangulation of manifolds. *Foundations of Computational Mathematics* (2017), 1–33.
- [3] BOISSONNAT, J.-D., DYER, R., GHOSH, A., AND OUDOT, S. Equating the witness and restricted Delaunay complexes. Tech. Rep. CGL-TR-24, Computational Geometric Learning, 2011. <http://cgl.uni-jena.de/Publications/WebHome>.
- [4] BOISSONNAT, J.-D., DYER, R., GHOSH, A., AND OUDOT, S. Y. Only distances are required to reconstruct submanifolds. Research report, INRIA Sophia Antipolis, 2014.
- [5] BOISSONNAT, J.-D., DYER, R., GHOSH, A., AND OUDOT, S. Y. Only distances are required to reconstruct submanifolds. *Comp. Geom. Theory and Appl.* (2016). To appear.
- [6] BOISSONNAT, J.-D., WORMSER, C., AND YVINEC, M. Anisotropic Delaunay mesh generation. *SIAM Journal on Computing* 44, 2 (2015), 467–512.
- [7] CAMPEN, M., HEISTERMANN, M., AND KOBELT, L. Practical anisotropic geodesy. In *Proceedings of the Eleventh Eurographics/ACMSIGGRAPH Symposium on Geometry Processing* (2013), SGP ’13, Eurographics Association, pp. 63–71.
- [8] CAÑAS, G. D., AND GORTLER, S. J. Orphan-free anisotropic Voronoi diagrams. *Discrete and Computational Geometry* 46, 3 (2011).
- [9] CAÑAS, G. D., AND GORTLER, S. J. Duals of orphan-free anisotropic Voronoi diagrams are embedded meshes. In *SoCG* (2012), ACM, pp. 219–228.
- [10] CAO, T., EDELSBRUNNER, H., AND TAN, T. Proof of correctness of the digital Delaunay triangulation algorithm. *Comp. Geo.: Theory and Applications* 48 (2015).
- [11] CHENG, S.-W., DEY, T. K., RAMOS, E. A., AND WENGER, R. Anisotropic surface meshing. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms* (2006), Society for Industrial and Applied Mathematics, pp. 202–211.
- [12] D’AZEVEDO, E. F., AND SIMPSON, R. B. On optimal interpolation triangle incidences. *SIAM J. Sci. Statist. Comput.* 10, 6 (1989), 1063–1075.
- [13] DEY, T. K., FAN, F., AND WANG, Y. Graph induced complex on point data. *Computational Geometry* 48, 8 (2015), 575–588.
- [14] DU, Q., AND WANG, D. Anisotropic centroidal Voronoi tessellations and their applications. *SIAM Journal on Scientific Computing* 26, 3 (2005), 737–761.

- [15] DYER, R., VEGTER, G., AND WINTRAECKEN, M. Riemannian simplices and triangulations. Preprint: arXiv:1406.3740.
- [16] DYER, R., ZHANG, H., AND MÖLLER, T. Surface sampling and the intrinsic Voronoi diagram. *Computer Graphics Forum* 27, 5 (2008), 1393–1402.
- [17] FUNKE, S., KLEIN, C., MEHLHORN, K., AND SCHMITT, S. Controlled perturbation for delaunay triangulations. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms* (2005), Society for Industrial and Applied Mathematics, pp. 1047–1056.
- [18] GARLAND, M., AND HECKBERT, P. S. Surface simplification using quadric error metrics. In *ACM SIGGRAPH* (1997), pp. 209–216.
- [19] KARCHER, H. Riemannian center of mass and mollifier smoothing. *Communications on Pure and Applied Mathematics* 30 (1977), 509–541.
- [20] LABELLE, F., AND SHEWCHUK, J. R. Anisotropic Voronoi diagrams and guaranteed-quality anisotropic mesh generation. In *SCG'03 : Proceedings of the Nineteenth Annual Symposium on Computational Geometry* (2003), ACM, pp. 191–200.
- [21] LEIBON, G. *Random Delaunay triangulations, the Thurston-Andreev theorem, and metric uniformization*. PhD thesis, UCSD, 1999.
- [22] MIREBEAU, J.-M. Optimal meshes for finite elements of arbitrary order. *Constructive approximation* 32, 2 (2010), 339–383.
- [23] NIYOGI, P., SMALE, S., AND WEINBERGER, S. Finding the homology of submanifolds with high confidence from random samples. *Discrete & Comp. Geom.* 39, 1-3 (2008).
- [24] PEYRÉ, G., PÉCHAUD, M., KERIVEN, R., AND COHEN, L. D. Geodesic methods in computer vision and graphics. *Found. Trends. Comput. Graph. Vis.* (2010).
- [25] ROURKE, C., AND SANDERSON, B. *Introduction to piecewise-linear topology*. Springer Science & Business Media, 2012.
- [26] ROUXEL-LABBÉ, M., WINTRAECKEN, M., AND BOISSONNAT, J.-D. Discretized Riemannian Delaunay triangulations. In *Proc. of the 25th Intern. Mesh. Round.* (2016), Elsevier.
- [27] SHEWCHUK, J. R. What is a good linear finite element? Interpolation, conditioning, anisotropy, and quality measures.
- [28] SPERNER, E. Fifty years of further development of a combinatorial lemma. *Numerical solution of highly nonlinear problems* (1980), 183–197.

Overview of the appendices

The topics of the appendices are as follows:

- A We review basic definitions related to simplices and complexes.
- B We describe our generalization of the notion of distortion between metrics due to Labelle and Shewchuk.
- C We discuss the separation of Voronoi objects. The main differences between this appendix and [2] are in our definition of metric distortion, the use of power protection, and the more geometrical nature of the proofs.
- D This appendix is related to the previous one and focuses on dihedral angles of Delaunay simplices. These results are intermediary steps used in Appendices E and H.
- E Here we built upon the previous two sections and discuss the stability of nets and Voronoi cells. This section distinguishes itself by the elementary and geometrical nature of the proofs.
- F We prove that the Delaunay simplices can be straightened, under sufficient conditions. The proof of the stability of the center of mass, which forms the core of this appendix, is also remarkable in the sense that it generalizes trivially to a far more general setting.
- G We illustrate a degenerate case of Section 6.3.
- H This appendix gives the proofs for our main result, Theorem 5.1, in the general setting of an arbitrary metric. We naturally rely heavily on Appendices C, D, E, and H.

The references for the Appendix can be found at the end of the Appendix.

A Simplices and complexes

The purpose of this section is to offer the precise definitions of concepts and notions related to simplicial complexes. The following definitions live within the context of abstract simplices and complexes.

A *simplex* σ is a non-empty finite set. The *dimension* of σ is given by $\dim \sigma = \#(\sigma) - 1$, and a j -simplex refers to a simplex of dimension j . The elements of σ are called the *vertices* of σ . The set of vertices of σ is noted $\text{Vert}(\sigma)$.

If a simplex τ is a subset of σ , we say it is a *face* of σ , and we write $\tau \leq \sigma$. A 1-dimensional face is called an *edge*. If τ is a proper subset of σ , we say it is a *proper* face and we write $\tau < \sigma$. A *facet* of σ is a face τ with $\dim \tau = \dim \sigma - 1$.

For any vertex $p \in \sigma$, the face opposite to p is the face determined by the other vertices of σ , and is denoted by σ_p . If τ is a j -simplex, and p is not a vertex of σ , we may construct a $(j+1)$ -simplex $\sigma = p * \tau$, called the *join* of p and τ as the simplex defined by p and the vertices of τ .

The *length* of an edge is the distance between its vertices. The *height* of p in σ is $D(p, \sigma) = d(p, \text{aff } \sigma_p)$.

A *circumscribing ball* for a simplex σ is any n -dimensional ball that contains the vertices of σ on its boundary. If σ admits a circumscribing ball, then it has a circumcenter, $C(\sigma)$, which is the center of the unique smallest circumscribing ball for σ . The radius of this ball is the *circumradius* of σ , denoted by $R(\sigma)$.

A.1 Complexes

Before defining Delaunay triangulations, we introduce the more general concept of simplicial complexes. Since the standard definition of a simplex as the convex hull of a set of points does not extend well to the Riemannian setting (see Dyer [15]), we approach these definitions from a more abstract point of view.

The *length* of an edge is the distance between its vertices. A *circumscribing ball* for a simplex σ is any n -dimensional ball that contains the vertices of σ on its boundary. If σ admits a circumscribing ball, then it has a circumcenter, $C(\sigma)$, which is the center of the unique smallest circumscribing ball for σ . The radius of this ball is the *circumradius* of σ , denoted by $R(\sigma)$. The *height* of p in σ is $D(p, \sigma) = d(p, \text{aff}(\sigma_p))$. The *dihedral angle* between two facets is the angle between their two supporting planes. If σ is a j -simplex with $j \geq 2$, then for any two vertices $p, q \in \sigma$, the dihedral angle between σ_p and σ_q defines an equality between ratios of heights (see Figure 6).

$$\sin \angle(\text{aff}(\sigma_p), \text{aff}(\sigma_q)) = \frac{D(p, \sigma)}{D(p, \sigma_q)} = \frac{D(q, \sigma)}{D(q, \sigma_p)}.$$




Figure 6: Acute and obtuse dihedral angles

A.2 Simplicial complexes

Simplicial complexes form the underlying framework of Delaunay triangulations. An *abstract simplicial complex* is a set \mathcal{K} of simplices such that if $\sigma \in \mathcal{K}$, then all the faces of σ also belong to \mathcal{K} . The union of the vertices of all the simplices of \mathcal{K} is called the *vertex set* of \mathcal{K} . The dimension of a complex is the largest dimension of any of its simplices. A subset $L \subseteq \mathcal{K}$ is a *subcomplex* of \mathcal{K} if it is also a complex. Two simplices are *adjacent* if they share a face and *incident* if one is a face of the other. If a simplex in \mathcal{K} is not a face of a larger simplex, we say that it is *maximal*. If all the maximal simplices in a complex \mathcal{K} of dimension n have dimension n , then the simplicial complex is *pure*. The *star* of a vertex p in a complex \mathcal{K} is the subcomplex S formed by set of simplices that are incident to p . The *link* of a vertex p is the union of the simplices opposite of p in S_p .

A *geometric simplicial complex* is an abstract simplicial complex whose faces are materialized, and to which the following condition is added: the intersection of any two faces of the complex is either empty or a face of the complex.

B Geodesic distortion

The concept of distortion was originally introduced by Labelle and Shewchuk [20] to relate distances with respect to two metrics, but this result can be (locally) extended to geodesic distances.

B.1 Original distortion

We first recall their definition and then show how to extend it to metric fields. The notion of metric transformation is required to define this original distortion, and we thus recall it now.

B.1.1 Metric transformation

Given a symmetric positive definite matrix G , we denote by F any matrix such that $\det(F) > 0$ and $F^t F = G$. The matrix F is called a *square root* of G . The square root of a matrix is not uniquely defined: the Cholesky decomposition provides, for example, an upper triangular F , but it is also possible to consider the diagonalization of G as $O^T D O$, where O is an orthonormal matrix and D is a diagonal matrix; the square root is then taken as $F = O^T \sqrt{D} O$. The latter definition is more natural than other decompositions since \sqrt{D} is canonically defined and F is then also a symmetric definite positive matrix, with the same eigenvectors as G . Regardless of the method chosen to compute the square root of a metric, we assume that it is fixed and identical for all the metrics considered.

The square root F offers a linear bijective transformation between the Euclidean space and a metric space, noting that

$$d_G(x, y) = \sqrt{(x - y)^t F^t F (x - y)} = \|F(x - y)\| = \|Fx - Fy\|,$$

where $\|\cdot\|$ stands for the Euclidean norm. Thus, the metric distance between x and y in Euclidean space can be seen as the Euclidean distance between two *transformed* points Fx and Fy living in the metric space of G .

B.1.2 Distortion

The *distortion* between two points p and q of Ω is defined by Labelle and Shewchuk [20] as $\psi(p, q) = \psi(G_p, G_q) = \max \{\|F_p F_q^{-1}\|, \|F_q F_p^{-1}\|\}$, where $\|\cdot\|$ is the Euclidean matrix norm, that is $\|M\| = \sup_{x \in \mathbb{R}^n} \frac{\|Mx\|}{\|x\|}$. Observe that $\psi(G_p, G_q) \geq 1$ and $\psi(G_p, G_q) = 1$ when $G_p = G_q$.

A fundamental property of the distortion is to relate the two distances d_{G_p} and d_{G_q} . Specifically, for any pair x, y of points, we have

$$\frac{1}{\psi(p, q)} d_{G_p}(x, y) \leq d_{G_q}(x, y) \leq \psi(p, q) d_{G_p}(x, y). \quad (3)$$

Indeed,

$$d_p(x, y) = \|F_p(x - y)\| = \|F_p F_q^{-1} F_q(x - y)\| \leq \|F_p F_q^{-1}\| \|F_q(x - y)\| \leq \psi(G_p, G_q) d_q(x, y),$$

and the other inequality is obtained similarly.

B.2 Geodesic distortion

The previous definition can be defined to hold (locally) for metric fields instead of metric, as we show now.

Lemma B.1. *Let $U \subset \Omega$ be open, and g and g' be two Riemannian metric fields on U . Let $\psi_0 \geq 1$ be a bound on the metric distortion, in the sense of Labelle and Shewchuk. Suppose that U is included in a ball $B_g(p_0, r_0)$, with $p_0 \in U$ and $r_0 \in \mathbb{R}^+$, such that $\forall p \in B(p_0, r_0)$, $\psi(g(p), g'(p)) \leq \psi_0$. Then, for all $x, y \in U$,*

$$\frac{1}{\psi_0} d_g(x, y) \leq d_{g'}(x, y) \leq \psi_0 d_g(x, y),$$

where d_g and $d_{g'}$ indicate the geodesic distances with respect to g and g' respectively.

Proof. Recall that for $p \in B_g(p_0, r_0)$ and, for any pair x, y of points, we have

$$\frac{1}{\psi_0} d_{g(p)}(x, y) \leq d_{g'(p)}(x, y) \leq \psi_0 d_{g(p)}(x, y).$$

Therefore, for any curve $\gamma(t)$ in U , we have that

$$\frac{1}{\psi_0} \int \sqrt{\langle \dot{\gamma}, \dot{\gamma} \rangle_{g(\gamma(t))}} dt \leq \int \sqrt{\langle \dot{\gamma}, \dot{\gamma} \rangle_{g'(\gamma(t))}} dt \leq \psi_0 \int \sqrt{\langle \dot{\gamma}, \dot{\gamma} \rangle_{g(\gamma(t))}} dt.$$

Considering the infimum over all paths γ that begin at x and end at y , we obtain the result. \square

Note that this result is independent from the definition of the distortion and is entirely based on the inequality comparing distances in two metrics (Equation 3).

C Separation of Voronoi objects

Power protected point sets were introduced to create quality bounds for the simplices of Delaunay triangulations built using such point sets [2]. We will show that power protection allows to deduce additional useful results for Voronoi diagrams. In this section we show that when a Voronoi diagram is built using a power protected sample set, its non-adjacent Voronoi faces, and specifically its Voronoi vertices are separated. This result is essential to our proofs in Sections 6 and 7 where we approximate complicated Voronoi cells with simpler Voronoi cells without creating inversions in the dual, which is only possible because we know that Voronoi vertices are far from one another.

We assume that the protection parameter δ is proportional to the sampling parameter ε , thus there exists a positive ι , with $\iota \leq 1$, such that $\delta = \iota\varepsilon$. We assume that the separation parameter μ is proportional to the sampling parameter ε and thus there exists a positive λ , with $\lambda \leq 1$, such that $\mu = \lambda\varepsilon$.

C.1 Separation of Voronoi vertices

The main result provides a bound on the Euclidean distance between Voronoi vertices of the Euclidean Voronoi diagram of a point set and is given by Lemma C.4. The following three lemmas are intermediary results needed to prove Lemma C.4.

Lemma C.1. *Let $B = B(c, R)$ and $B' = B(c', R')$ be two n -balls whose bounding spheres ∂B and $\partial B'$ intersect, and let H be the bisecting hyperplane of B and B' , i.e. the hyperplane that*

contains the $(n-2)$ -sphere $S = \partial B \cap \partial B'$. Let θ be the angle of the cone (c, S) . Writing $\rho = \frac{R'}{R}$ and $\|c - c'\| = \lambda R$, we have

$$\cos(\theta) = \frac{1 + \lambda^2 - \rho^2}{2\lambda}. \quad (4)$$

If $R \geq R'$, we have $\cos(\theta) \geq \frac{\lambda}{2}$.

Proof. Let $q \in S$; applying the cosine rule to the triangle $\triangle cc'q$ gives

$$\lambda^2 R^2 + R^2 - 2\lambda R^2 \cos(\theta) = R'^2, \quad (5)$$

which proves Equation (4). If $R \geq R'$, then $\rho \leq 1$, and $\cos(\theta) \geq \lambda/2$ immediately follows from Equation (4). \square




Figure 7: Construction used in Lemmas C.1 and C.2.

Lemma C.2. Let $B = B(c, R)$ and $B' = B(c', R')$ be two n -balls whose bounding spheres ∂B and $\partial B'$ intersect, and let $\tilde{\theta}$ be the angle of the cone (c, \tilde{S}) where $\tilde{S} = \partial B \cap \partial B'^{\delta}$. Writing $\|c - c'\| = \lambda R$, we have

$$\cos(\tilde{\theta}) = \cos(\theta) - \frac{\delta^2}{2R^2\lambda}$$

Proof. Let $\tilde{q} \in \tilde{S}$, applying the cosine rule to the triangle $\triangle cc'\tilde{q}$ gives

$$\lambda^2 R^2 + R^2 - 2\lambda R^2 \cos(\tilde{\theta}) = R'^2 + \delta^2.$$

Subtracting (5) from the previous equality yields $\delta^2 = 2\lambda R^2(\cos(\theta) - \cos(\tilde{\theta}))$, which proves the lemma. \square

The altitude of the vertex p_i in the simplex σ is denoted by $D(p_i, \sigma)$.

Lemma C.3. *Let $\sigma = p * \tau$ and $\sigma' = p' * \tau$ be two Delaunay simplices sharing a common facet τ . (Here $*$ denotes the join operator.) Let $B(\sigma) = B(c, R)$ and $B(\sigma') = B(c', R')$ be the circumscribing balls of σ and σ' respectively. Then σ' is δ -power protected with respect to p , that is $p \notin B(\sigma')^{+\delta}$ if and only if $\|c - c'\| \geq \frac{\delta^2}{2D(p, \sigma)}$.*

Proof. The spheres ∂B and $\partial B'^{+\delta}$ intersect in a $(n-2)$ -sphere \tilde{S} which is contained in a hyperplane \tilde{H} parallel to the hyperplane $H = \text{aff}(\tau)$. For any $\tilde{q} \in \tilde{S}$ we have

$$d(\tilde{H}, H) = d(\tilde{q}, H) = R(\cos(\theta) - \cos(\tilde{\theta})) = \frac{\delta^2}{2\|c - c'\|},$$

where the last equality follows from Lemma C.2 and $d(\tilde{H}, H)$ denotes the distance between the two parallel hyperplanes. See Figure 7 for a sketch. Since $p \in \partial B$, p belongs to $B(\sigma')^{+\delta}$ if and only if p lies in the strip bounded by H and \tilde{H} , which is equivalent to

$$d(p, H) = D(p, \sigma) < \frac{\delta^2}{2\|c - c'\|}.$$

The result now follows. \square

We can make this bound independent of the simplices considered, as shown in Lemma C.4.

Lemma C.4. *Let \mathcal{P} be a δ -power protected (ε, μ) -sample. The protection parameter ι is given by $\delta = \iota\varepsilon$. For any two adjacent Voronoi vertices c and c' of $\text{VD}(\mathcal{P})$, we have*

$$\|c - c'\| \geq \frac{\delta^2}{4\varepsilon} = \frac{\iota^2\varepsilon}{4}.$$

Proof. For any simplex σ , we have $D(p, \sigma) \leq 2R_\sigma$ for all $p \in \sigma$, where R_σ denotes the radius of the circumsphere of σ . For any σ in the triangulation of an ε -net, we have $R_\sigma \leq \varepsilon$. Thus $D(p, \sigma) \leq 2\varepsilon$, and Lemma C.3 yields $\|c - c'\| \geq \delta^2/4\varepsilon$. \square

Remark C.5. *In this section, we have computed a lower bound on the distance between two (adjacent) Voronoi vertices. In Appendix E, we shall show that Voronoi vertices are stable under metric perturbations, meaning that when a metric field is slightly modified, the position of a Voronoi vertex does not move too much. The combination of this separation and stability will then be the basis of many proofs in this paper.*

C.2 Separation of Voronoi faces (Proof of Lemma 6.5)

Similar results can be obtained on the distance between a Voronoi vertex c and faces that are not incident to c , also referred to as *foreign* faces. Note that we are still in the context of an Euclidean metric.

The following lemma can be found in [3, Lemma 3.3] for ordinary protection instead of power protection.

Lemma C.6. *Suppose that c is the circumcenter of a δ -power protected simplex σ of a Delaunay triangulation built from an ε -sample, then all foreign Voronoi faces are at least $\delta^2/8\varepsilon$ removed from c .*




Figure 8: Illustration of the notations for the proof of Lemma 6.5. The simplex σ is dashed in yellow and has vertices $p_0 p_1 p_2$. The distances cc' and cc'' are lower bounded by $\delta^2/4\varepsilon$.

Proof. We denote by p_0 an (arbitrary) vertex of σ , and by q a vertex that is not in σ but is adjacent to p_0 in $\text{Del}(\mathcal{P})$. Let x be a point in $B(c, r) \cap V_{\mathbb{E}}(p_0)$, with $0 < r < \delta^2/4\varepsilon$. The upper bound for r is chosen with Lemma C.3 in mind: we are trying to find a condition such that $x \in V_{\mathbb{E}}(p_0)$. The notations are illustrated in Figure 8.

Because of the triangle inequality, we have that

$$\begin{aligned} |d(c, q) - d(x, q)| &\leq d(x, c) \\ |d(c, p_i) - d(x, p_i)| &\leq d(x, c). \end{aligned}$$

By power protection, we have that $d(c, q)^2 \geq d(c, p_i)^2 + \delta^2$. Therefore,

$$\begin{aligned} (d(x, q) + r)^2 &\geq (d(x, p_i) - r)^2 + \delta^2 \\ d(x, q)^2 + 2rd(x, q) &\geq d(x, p_i)^2 - 2rd(x, p_i) + \delta^2 \\ d(x, q)^2 &\geq d(x, p_i)^2 - 2r(d(x, p_i) + d(x, q)) + \delta^2. \end{aligned}$$

Without loss of generality, we can assume that q is the site closest to c and thus $d(x, q) < d(x, c)$. If \mathcal{P} is an ε -net, we have

$$d(x, p_i) + d(x, q) \leq d(x, p_i) + d(c, q) \leq \varepsilon + 3\varepsilon = 4\varepsilon,$$

so

$$d(x, q)^2 \geq d(x, p_i)^2 - 8r\varepsilon + \delta^2.$$

This implies that as long $r < \delta^2/8\varepsilon$, x lies in a Voronoi object associated to the vertices p_i of σ . \square

Further progressing, we can show that Voronoi faces are thick, with Lemma C.7. This property is useful to construct a triangulation that satisfies the hypotheses of Sperner's lemma.

Lemma C.7. *Let \mathcal{P} be a δ -power protected (ε, μ) -net. Let V_0 be the Voronoi cell of the site $p_0 \in \mathcal{P}$ in the Euclidean Voronoi diagram $\text{VD}_{\mathbb{E}}(\mathcal{P})$. Then for any k -face F_0 of V_0 with $k \in [1, \dots, n]$, there exists $x \in F_0$ such that*

$$d(x, \partial F_0) > \frac{\delta^2}{16\varepsilon},$$

where ∂F_0 denotes the boundary of the face F_0 .

Proof. All the vertices of F_0 are circumcenters of $\text{VD}_{\mathbb{E}}(\mathcal{P})$. Consider the erosion of the face F_0 by a ball of radius $\frac{\delta^2}{16\varepsilon}$ and denote it F_0^- . If F_0^- is empty, we contradict the separation Lemma 6.5. \square

D Bounds on dihedral angles

The use of nets allows us to deduce bounds on the dihedral angles of faces of the Delaunay triangulation, as well as on the dihedral angles between adjacent faces of a Voronoi diagram. Those bounds are frequently used throughout this paper, and specifically to prove stability of Voronoi vertices (see Appendix E). Since we are interested in dihedral bounds in the Euclidean setting, the point set is first assumed to be a net with respect to the Euclidean metric field. We complicate matters slightly with Lemma D.5 by assuming that the point set is a power protected net with respect to an arbitrary metric field that is not too far from the Euclidean metric field (in terms of distortion), and still manage to expose bounds with respect to the Euclidean distance.

D.1 Bounds on the dihedral angles of Euclidean Voronoi cells

Assuming that a point set is an (ε, μ) -net allows us to deduce lower and upper bounds on the dihedral angles between adjacent Voronoi faces when the metric field is Euclidean.

Lemma D.1. *Let $\Omega = \mathbb{R}^n$ and \mathcal{P} be an (ε, μ) -net with respect to the Euclidean distance on Ω . Let $p \in \mathcal{P}$ and $V(p)$ be the Voronoi cell of $p \in \mathcal{P}$. Let $q, r \in \mathcal{P}$ be two sites such that $V(q)$ and $V(r)$ are adjacent to $V(p)$ in the Euclidean Voronoi diagram of \mathcal{P} . Let θ be the dihedral angle between $\text{BS}(p, q)$ and $\text{BS}(p, r)$. Then*

$$2 \arcsin\left(\frac{\mu}{2\varepsilon}\right) \leq \theta \leq \pi - \arcsin\left(\frac{\mu}{2\varepsilon}\right).$$

Proof. We consider the plane \mathcal{H} that passes through the sites p, q and r . Notations used below are illustrated in Figure 9.

Lower bound. Let m_{pq} and m_{pr} be the projections of the site p on respectively the bisectors $\text{BS}(p, q)$ and $\text{BS}(p, r)$. Since \mathcal{P} is an (ε, μ) -net, we have that $l_q = |p - m_{pq}| \geq \mu/2$, $l_r = |p - m_{pr}| \geq \mu/2$ and $L = |p - c| \leq \varepsilon$. Thus

$$\theta = \arcsin\left(\frac{l_q}{L}\right) + \arcsin\left(\frac{l_r}{L}\right) \geq 2 \arcsin\left(\frac{\mu/2}{\varepsilon}\right) = 2 \arcsin\left(\frac{\mu}{2\varepsilon}\right).$$

Note that since $0 < \mu < 2\varepsilon$, we have $0 < \mu/2\varepsilon < 1$.

Upper bound. To obtain an upper bound on θ , we compute a lower bound on the angle $\alpha = \widehat{qpr}$ at p , noting that $\theta = \pi - \alpha$. Let $l_{qr} = |q - r|$, and $R = |c - r|$. By the law of sines, we have

$$\frac{l_{qr}}{\sin(\alpha)} = 2R.$$




Figure 9: Construction and notations used in Lemma D.1.

Since \mathcal{P} is an (ε, μ) -net, we have $l_{qr} \geq \mu$ and $R \leq \varepsilon$. Finally,

$$\alpha \geq \arcsin\left(\frac{\mu}{2\varepsilon}\right) \implies \theta \leq \pi - \arcsin\left(\frac{\mu}{2\varepsilon}\right).$$

□

D.2 Bounds on the dihedral angles of Euclidean Delaunay simplices

Bounds on the dihedral angles of simplices guarantee the thickness – the smallest height of any vertex – of simplices, and thus their quality. Additionally, they can be used to show that circumcenters of adjacent simplices are far from one another, thus proving the stability of circumcenters and of Delaunay simplices.

D.2.1 Using power protection with respect to the Euclidean metric field

We first assume that the metric field g is the Euclidean metric field g_E and show that the simplices of an Euclidean Delaunay triangulation constructed from a power protected net are thick.

Recall that the dihedral angle can be expressed through heights as

$$\sin \angle(\text{aff}(\sigma_p), \text{aff}(\sigma_q)) = \frac{D(p, \sigma)}{D(p, \sigma_q)} = \frac{D(q, \sigma)}{D(q, \sigma_p)}.$$

The bound on dihedral angles is obtained by bounding the height of vertices in a simplex. An obvious upper bound on the height of a vertex p in σ is $D(p, \sigma) < 2\varepsilon$. A lower bound is already obtained in Lemma C.3: we have that $D(p, \sigma) \geq \delta^2/2 \|c - c'\| = \delta^2/4\varepsilon$. We can thus bound the dihedral angles as follows:

Lemma D.2. *Let \mathcal{P} be a δ -power protected (ε, μ) -net with respect to the Euclidean metric field g_0 . Let φ be the dihedral angle between two facets τ_1 and τ_2 of a simplex σ of $\text{Del}_{g_0}(\mathcal{P})$. Then*

$$\arcsin(s_0) \leq \varphi \leq \pi - \arcsin(s_0),$$

with $s_0 = \frac{A_{\lambda, \varepsilon}}{2}$ and $A_{\lambda, \varepsilon}$ defined as in the previous Lemma.

Proof. Recall that

$$\sin(\varphi) = \sin \angle(\text{aff}(\sigma_p), \text{aff}(\sigma_q)) = \frac{D(p, \sigma)}{D(p, \sigma_q)} = \frac{D(q, \sigma)}{D(q, \sigma_p)}.$$

From previous remarks, we have that

$$D(q, \sigma_p) \geq D(q, \sigma) \geq \frac{\delta^2}{4\varepsilon},$$

and $D(q, \sigma) \leq 2\varepsilon$. Thus, if $\varphi = \angle(\text{aff}(\sigma_p), \text{aff}(\sigma_q))$, then

$$\sin \angle(\text{aff}(\sigma_p), \text{aff}(\sigma_q)) \geq \frac{\delta^2}{4\varepsilon} \frac{1}{2\varepsilon} = \frac{\iota^2}{2} =: s_0$$

Note that $0 < s_0 < 1$ and thus

$$\arcsin(s_0) \leq \varphi \leq \pi - \arcsin(s_0).$$

□

D.2.2 Using power protection with respect to an arbitrary metric field

When considering a Voronoi diagram built using the geodesic distance induced by an arbitrary metric field g , the assumption of a power protected net is made with respect to this geodesic distance. To prove the stability of the power protected assumption under metric perturbation, we will however need to deduce lower and upper bounds on the dihedral angles between faces of the simplices of the Riemannian Delaunay complex with respect to the Euclidean metric field. We prove here that if the point set \mathcal{P} is a δ -power protected (ε, μ) -net with respect to an arbitrary metric field g and if the distortion between g and the Euclidean metric field $g_{\mathbb{E}}$ is small, then the dihedral angles of the simplices of the Euclidean Delaunay triangulation of \mathcal{P} can be bounded.

We first give a result on the stability of Delaunay balls which expresses that if two metric fields have low distortion, the Delaunay balls of a simplex with respect to each metric field are close. One of these metric fields is assumed to be the Euclidean metric field. A similar result can be found in the proof of Lemma 5 in the theoretical analysis of locally uniform anisotropic meshes of Boissonnat et al. [6].

Lemma D.3. *Let $U \subset \Omega$ be open, and g and g' be two Riemannian metric fields on U . Let $\psi_0 \geq 1$ be a bound on the metric distortion. Suppose that U is included in a ball $B_g(p_0, r_0)$, with $p_0 \in U$ and $r_0 \in \mathbb{R}^+$, such that $\forall p \in B(p_0, r_0), \psi(g(p), g'(p)) \leq \psi_0$. Assume furthermore that g' is the Euclidean distance (thus $d_{g'} = d_{\mathbb{E}}$). Let $B = B_g(c, r)$ be the geodesic ball with respect to the metric field g , centered on $c \in U$ and of radius r . Assume that $B_{\mathbb{E}}(c, \psi_0 r) \subset U$. Then B can be encompassed by two Euclidean balls $B_{\mathbb{E}}(c, r_{-\psi_0})$ and $B_{\mathbb{E}}(c, r_{+\psi_0})$ with $r_{-\psi_0} = r/\psi_0$ and $r_{+\psi_0} = \psi_0 r$.*

Proof. This is a straight consequence from Lemma B.1. Indeed, we have for all $x \in U$ that

$$\frac{1}{\psi_0} d_{\mathbb{E}}(c, x) \leq d_g(c, x) \leq \psi_0 d_{\mathbb{E}}(c, x),$$

and similarly

$$\frac{1}{\psi_0} d_g(c, x) \leq d_{\mathbb{E}}(c, x) \leq \psi_0 d_g(c, x).$$

Thus,

$$x \in B_{\mathbb{E}}\left(c, \frac{r}{\psi_0}\right) \iff d_{\mathbb{E}}(c, x) \leq \frac{r}{\psi_0} \implies \frac{1}{\psi_0} d_g(c, x) \leq \frac{r}{\psi_0},$$

giving us $B_{\mathbb{E}}(c, r_{-\psi_0}) \subset B$. On the other hand, we have

$$x \in B_g(c, r) \iff d_g(c, x) \leq r \implies \frac{1}{\psi_0} d_{\mathbb{E}}(c, x) \leq r,$$

giving us $B \subset B_{\mathbb{E}}(c, r_{+\psi_0})$. \square

Note that $r_{-\psi_0}$ and $r_{+\psi_0}$ go to r as ψ_0 goes to 1.

We now use this stability result to provide Euclidean dihedral angle bounds assuming power protection with an arbitrary metric field that is close to $g_{\mathbb{E}}$. We first require the intermediary result given by Lemma D.4.

Lemma D.4 (Whitney's lemma). *Let H be a hyperplane in Euclidean n -space and τ an $n-1$ -simplex whose vertices lie at most η from the H and whose minimum height is h_{\min} . Then the angle ξ between $\text{aff}(\tau)$ and H is bounded from above by*

$$\sin(\xi) \leq \frac{\eta d}{h_{\min}}.$$

Proof. By definition, the barycenter of a $(n-1)$ -simplex has barycentric coordinates $\lambda_i = 1/d$. This means that it has distance h_{\min}/n to each of its faces. So the ball centered on the barycenter with radius h_{\min}/n is contained in the simplex. This means that for any direction in $\text{aff}(\tau)$ there exists a line segment of length $2h_{\min}/n$ that lies within τ . Moreover the end points of this line segments lie at most η from H because the vertices of the τ do. This means that the angle ξ is bounded by

$$\sin(\xi) \leq \frac{2\eta}{2h_{\min}/d} = \frac{\eta d}{h_{\min}}.$$

\square


We can now give the main result which bounds *Euclidean* dihedral angles, assuming power protection with respect to the arbitrary metric field.

Lemma D.5. *Let $U \subset \Omega$ be open, and g and g' be two Riemannian metric fields on U . Let $\psi_0 \geq 1$ be a bound on the metric distortion. Suppose that U is included in a ball $B_g(p_0, r_0)$, with $p_0 \in U$ and $r_0 \in \mathbb{R}^+$, such that $\forall p \in B(p_0, r_0), \psi(g(p), g'(p)) \leq \psi_0$. Let \mathcal{P}_U be a δ -power protected (ε, μ) -net over U , with respect to g . Let $B = B_g(c, r)$ and $B' = B_g(c', r')$ be the geodesic Delaunay balls of $\sigma = \tau * p$ and $\sigma' = \tau * p'$, with $\sigma, \sigma' \in \text{Del}_g(\mathcal{P}_U)$. Assume that \mathcal{P}_U is sufficiently dense such that U contains B and B' . Then the minimum height of the simplex satisfies*

$$h_{\min} = \sqrt{1 - \left(n \frac{(r^2 + r'^2) \left(\frac{1}{\psi_0^2} - \psi_0^2 \right)}{8\varepsilon h_{\min}} \right)^2} \cdot \left(\frac{(r^2 + r'^2) \left(\frac{1}{\psi_0^2} - \psi_0^2 \right) + \frac{\delta^2}{\psi_0^2}}{4\varepsilon} - \frac{n \frac{(r^2 + r'^2) \left(\frac{1}{\psi_0^2} - \psi_0^2 \right)}{8\varepsilon h_{\min}}}{\sqrt{1 - \left(n \frac{(r^2 + r'^2) \left(\frac{1}{\psi_0^2} - \psi_0^2 \right)}{8\varepsilon h_{\min}} \right)^2}} (r + r') \right).$$

We proceed in a similar fashion as the proof Lemma C.1. However, a significant difference is that we are interested here in only proving that power protection with respect to the generic metric field provides a height bound in the Euclidean setting (rather than an equivalence).

- $B_{\mathbb{E}}^{\pm\psi_0}$ and $B'_{\mathbb{E}}^{\pm\psi_0}$ are the two sets of (Euclidean) enclosing balls of respectively B and B' defined as in Lemma D.3.
- $B'^{+\delta}$ is the power protected ball of σ' , given by $B'^{+\delta} = B(c', \sqrt{r'^2 + \delta^2})$.
- $B_{\mathbb{E}}'^{+\delta, \pm\psi_0}$ are the two Euclidean balls enclosing $B'^{+\delta}$.
- $S = \partial B^{-\psi_0} \cap \partial B'^{+\psi_0}$, $\tilde{S} = \partial B^{+\psi_0} \cap \partial B'^{+\delta, -\psi_0}$ and $S' = \partial B^{+\psi_0} \cap \partial B'^{-\psi_0}$.
- q is a point on S , \tilde{q} is a point on \tilde{S} and q' is a point on S' .
- H is the geodesic supporting plane of τ , that is $\{\operatorname{argmin}(\sum_{v_i \in \tau} \lambda_i d_g(x, v_i))\}$.
- $H_{\mathbb{E}}$, $\tilde{H}_{\mathbb{E}}$ and $H'_{\mathbb{E}}$ are the two Euclidean hyperplanes orthogonal to $[cc']$ passing through respectively q , \tilde{q} and q' .
- $\theta = \widehat{c'cq}$, $\tilde{\theta} = \widehat{c'\tilde{c}\tilde{q}}$, and $\lambda_c = |cc'|/r$.



While the vertices of τ live on H , $\text{aff}_{\mathbb{E}}(\tau)$ is not necessarily orthogonal to $[cc']$ and consequently

$$d_{\mathbb{E}}(p, H_{\mathbb{E}}) \leq d_{\mathbb{E}}(p, \tau).$$

The separation between the hyperplanes $H_{\mathbb{E}}$ and $\tilde{H}_{\mathbb{E}}$ provides a lower bound on the distance $d_{\mathbb{E}}(p, \tau)$, thus on the height $D_{\mathbb{E}}(p, \sigma)$. We therefore seek to bound $d_{\mathbb{E}}(H_{\mathbb{E}}, \tilde{H}_{\mathbb{E}})$.

By definition of the enclosing Euclidean balls, we have that

$$|cq| = \frac{r}{\psi_0}, \quad |c\tilde{q}| = \psi_0 r, \quad |c'q| = \psi_0 r', \quad \text{and} \quad |c'\tilde{q}| = \frac{1}{\psi_0} \sqrt{r'^2 + \delta^2}$$

Using the law of cosines in the triangles $\triangle cc'q$ and $\triangle cc'\tilde{q}$, we find

$$\begin{aligned} \frac{r'^2 + \delta^2}{\psi_0^2} &= \lambda_c^2 r^2 + \psi_0^2 r'^2 - 2\lambda_c \psi_0 r'^2 \cos(\tilde{\theta}) \\ \psi_0^2 r'^2 &= \lambda_c^2 r^2 + \frac{r^2}{\psi_0^2} - 2\frac{\lambda_c r^2}{\psi_0} \cos(\theta), \end{aligned}$$

where $\lambda_c r = |c - c'|$. Subtracting one from the other, we obtain

$$\begin{aligned} \psi_0^2 r'^2 - \frac{r'^2 + \delta^2}{\psi_0^2} &= \frac{r^2}{\psi_0^2} - \psi_0^2 r'^2 + 2\lambda_c \psi_0 r'^2 \cos(\tilde{\theta}) - 2\frac{\lambda_c r^2}{\psi_0} \cos(\theta) \\ \iff \frac{r}{\psi_0} \cos(\theta) - \psi_0 r \cos(\tilde{\theta}) &= \frac{\frac{r^2}{\psi_0^2} - \psi_0^2 r'^2 + \frac{r'^2 + \delta^2}{\psi_0^2} - \psi_0^2 r'^2}{2\lambda_c r} \\ \iff \frac{r}{\psi_0} \cos(\theta) - \psi_0 r \cos(\tilde{\theta}) &= \frac{(r^2 + r'^2) \left(\frac{1}{\psi_0^2} - \psi_0^2 \right) + \frac{\delta^2}{\psi_0^2}}{2\lambda_c r}, \end{aligned}$$

so that

$$d(H_{\mathbb{E}}, \tilde{H}_{\mathbb{E}}) = \frac{r}{\psi_0} \cos(\theta) - \psi_0 r \cos(\tilde{\theta}) = \frac{(r^2 + r'^2) \left(\frac{1}{\psi_0^2} - \psi_0^2 \right) + \frac{\delta^2}{\psi_0^2}}{2\lambda_c r}.$$

Similarly we can calculate the distance $d_{\mathbb{E}}(H_{\mathbb{E}}, H'_{\mathbb{E}})$ to be:

$$\begin{aligned} d_{\mathbb{E}}(H_{\mathbb{E}}, H'_{\mathbb{E}}) &= \frac{\frac{1}{\psi^2}(r')^2 + |c - c'|^2 - \psi^2 r^2}{2|c - c'|} - \frac{\psi^2 (r')^2 + |c - c'|^2 - \frac{1}{\psi^2} r^2}{2|c - c'|} \\ &= \frac{(r^2 + r'^2) \left(\frac{1}{\psi_0^2} - \psi_0^2 \right)}{2|c - c'|} \\ &\leq \frac{(r^2 + r'^2) \left(\frac{1}{\psi_0^2} - \psi_0^2 \right)}{4\varepsilon} \end{aligned}$$

Lemma D.4 gives us that the angle ξ between $H_{\mathbb{E}}$ and $\text{aff}(\tau)$ is bounded by

$$\sin(\xi) \leq \frac{n d_{\mathbb{E}}(H_{\mathbb{E}}, H'_{\mathbb{E}})}{2h_{\min}}.$$

The vertices of τ lie in between $H_{\mathbb{E}}$ and $H'_{\mathbb{E}}$ and inside $B_{\mathbb{E}}(c', \psi r')$. Thus, if we restrict to the $\tilde{q}cc'$ plane, distance between the point where $\text{aff}(\tau)$ intersects $H_{\mathbb{E}}$ and the orthogonal projection $\pi_{H_{\mathbb{E}}}(\tilde{q})$ of \tilde{q} on $H_{\mathbb{E}}$ is at most $\psi(r + r')$. This in turn implies that the line connecting $\pi_{H_{\mathbb{E}}}(\tilde{q})$ and \tilde{q} intersects $\text{aff}(\tau)$ at most distance $(r + r') \tan(\xi)$ from $H_{\mathbb{E}}$. This also gives us that the distance

from \tilde{q} to its orthogonal projection $\pi_{\text{aff}(\tau)}(\tilde{q})$ on $\text{aff}(\tau)$ is bounded by

$$\cos(\xi)(d(\tilde{H}_{\mathbb{E}}, H_{\mathbb{E}}) - \tan(\xi)(r + r')) = \sqrt{1 - \left(\frac{n d_{\mathbb{E}}(H_{\mathbb{E}}, H'_{\mathbb{E}})}{2h_{\min}} \right)^2} \cdot \left(d(\tilde{H}_{\mathbb{E}}, H_{\mathbb{E}}) - \frac{\frac{n d_{\mathbb{E}}(H_{\mathbb{E}}, H'_{\mathbb{E}})}{2h_{\min}}}{\sqrt{1 - \left(\frac{n d_{\mathbb{E}}(H_{\mathbb{E}}, H'_{\mathbb{E}})}{2h_{\min}} \right)^2}} (r + r') \right).$$

Here we note that h_{\min} originally referred to the minimum height of a face, but because the height of a face is always greater than the height in the simplex we may read this in a general way, that is we regard h_{\min} as a universal lower bound on the height. Because $|\tilde{q} - \pi_{\text{aff}(\tau)}(\tilde{q})|$ bounds the height of the simplex we get the following relation:

$$\begin{aligned} h_{\min} &= \sqrt{1 - \left(\frac{n d_{\mathbb{E}}(H_{\mathbb{E}}, H'_{\mathbb{E}})}{2h_{\min}} \right)^2} \left(d(\tilde{H}_{\mathbb{E}}, H_{\mathbb{E}}) - \frac{\frac{n d_{\mathbb{E}}(H_{\mathbb{E}}, H'_{\mathbb{E}})}{2h_{\min}}}{\sqrt{1 - \left(\frac{n d_{\mathbb{E}}(H_{\mathbb{E}}, H'_{\mathbb{E}})}{2h_{\min}} \right)^2}} (r + r') \right) \\ h_{\min} &= \sqrt{1 - \left(n \frac{(r^2 + r'^2) \left(\frac{1}{\psi_0^2} - \psi_0^2 \right)}{8\varepsilon h_{\min}} \right)^2} \cdot \left(\frac{(r^2 + r'^2) \left(\frac{1}{\psi_0^2} - \psi_0^2 \right) + \frac{\delta^2}{\psi_0^2}}{4\varepsilon} - \frac{n \frac{(r^2 + r'^2) \left(\frac{1}{\psi_0^2} - \psi_0^2 \right)}{8\varepsilon h_{\min}}}{\sqrt{1 - \left(n \frac{(r^2 + r'^2) \left(\frac{1}{\psi_0^2} - \psi_0^2 \right)}{8\varepsilon h_{\min}} \right)^2}} (r + r') \right) \end{aligned}$$

To make the expression a bit more readable, we introduce

$$s_1 = \frac{(r^2 + r'^2) \left(\frac{1}{\psi_0^2} - \psi_0^2 \right)}{4\varepsilon}$$

so that

$$\begin{aligned} h_{\min} &= \sqrt{1 - \left(n \frac{s_1}{2h_{\min}} \right)^2} \left(\left(s_1 + \frac{\delta^2}{4\varepsilon\psi_0^2} \right) - \frac{n \frac{s_1}{2h_{\min}}}{\sqrt{1 - \left(n \frac{s_1}{2h_{\min}} \right)^2}} (r + r') \right) \\ h_{\min} + n \frac{s_1}{2h_{\min}} (r + r') &= \sqrt{1 - \left(n \frac{s_1}{2h_{\min}} \right)^2} \left(s_1 + \frac{\delta^2}{4\varepsilon\psi_0^2} \right). \end{aligned}$$

Multiplying with h_{\min} and squaring we find:

$$(h_{\min}^2 + n \frac{s_1}{2} (r + r'))^2 = \left(h_{\min}^2 - \left(n \frac{s_1}{2} \right)^2 \right) \left(s_1 + \frac{\delta^2}{4\varepsilon\psi_0^2} \right)^2$$

We note that in the limit $\psi \rightarrow 1$, s_1 tends to zero, we therefore expand h_{\min} in terms of s_1 . Using a computer algebra system we find

$$h_{\min}^2 = \frac{\delta^4}{2^{10}\psi^4\varepsilon^2} + \left(\frac{\delta^2}{128\psi^2\varepsilon} - n(r+r') \right) s_1 - \frac{\delta^4(16n^2-1) + 2^{14}n^2\psi^4\varepsilon^2(r+r')^2}{64\delta^4} s_1^2 + \mathcal{O}(s_1^3)$$

We emphasize that this equation gives $h_{\min} = \delta^2/4\varepsilon$ as ψ tends to 1. This means that for sufficiently small metric distortion the height of a protected simplex will be strictly positive. \square

Lemma D.6. *Let g be a metric field and \mathcal{P} be a point set defined over \mathbb{R}^n . Let $\psi_0 = \psi(g, \mathbb{E})$. Assume that \mathcal{P} is a δ -power protected (ε, μ) -net over \mathbb{R}^n . Let ϕ be the dihedral angle between two faces of a simplex $\tau \in \text{Del}_{\mathbb{E}}(\mathcal{P})$. Then*

$$\arcsin(s_0) \leq \phi \leq \pi - \arcsin(s_0),$$

with s_0 detailed in the proof.

Proof. Denote h_{\min} the lower bound on $D(q, \sigma)$ obtained in the previous lemma. We also immediately have that

$$D(q, \sigma) \leq 2\varepsilon.$$

Let φ be the dihedral angle between $\text{aff}(\sigma_p)$ and $\text{aff}(\sigma_q)$. Recall that

$$\sin(\varphi) = \sin \angle(\text{aff}(\sigma_p), \text{aff}(\sigma_q)) = \frac{D(p, \sigma)}{D(p, \sigma_q)} = \frac{D(q, \sigma)}{D(q, \sigma_p)}$$

Thus

$$\sin(\varphi) \geq \frac{h_{\min}}{2\varepsilon} =: s_0.$$

For s_0 to make sense, we want $s_0 > 0$, which is bound to happen as ψ_0 goes to 1, as shown in Lemma D.5: h_{\min} goes to $\delta^2/4\varepsilon$ thus $h_{\min}/4\varepsilon$ goes to $\iota^2/4 > 0$. \square

E Stability

The notion of stability designates the conservation of a property despite changes of other parameters. In our context, the main assumptions concern the nature of point sets: we assume that point sets are power protected nets and wish to preserve these hypotheses despite (small) metric perturbations. The stability is important both from a theoretical and a practical point of view. Indeed, if an assumption is stable under perturbation, we can simplify matters without losing information. For example, we will prove that if a point set is a net with respect to a metric field g , then it is also a net (albeit with slightly different sampling and separation parameters) for a metric field g' that is close to g (in terms of distortion). In a practical context, the stability of an assumption provides robustness with respect to numerical errors (see, for example, the work of Funke et al. [17] on the stability of Delaunay simplices).

E.1 Stability of the protected net hypothesis under metric transformation

It is rather immediate that the power protected net property is preserved when the point set is transformed (see Section B.1.1) between these spaces, as shown by the following lemma.

Lemma E.1. *Let \mathcal{P} be a δ -power protected (ε, μ) -net in Ω . Let $g = g_0$ be a uniform Riemannian metric field and F_0 a square root of g_0 . If \mathcal{P} is a δ -power protected (ε, μ) -net with respect to g_0 then $\mathcal{P}' = \{F_0 p_i, p_i \in \mathcal{P}\}$ is a δ -power protected (ε, μ) -net with respect to the Euclidean metric.*

Proof. This results directly from the observation that

$$d_0(x, y)^2 = (x - y)^t g_0(x - y) = \|F_0(x - y)\|^2 = d(F_0 x, F_0 y)^2.$$

□

E.2 Stability of the protected net hypothesis under metric perturbation

Metric field perturbations are small modifications of a metric field in terms of distortion. Since a generic Riemannian metric field g is difficult to study, we will generally consider a uniform approximation g_0 of g within a small neighborhood, such that the distortion between both metric fields is small over that neighborhood. In that context, the perturbation of g is the act of bringing g onto g_0 . Stability of the assumption of power protection was previously investigated by Boissonnat et al. [4] in the context of manifold reconstructions.

E.2.1 Stability of the net property

The following lemma shows that the “net” property is preserved when the metric field is perturbed: a point set that is a net with g is also a net with respect to g' , a metric field that is close to g .

Lemma E.2. *Let $U \subset \Omega$ be open, and g and g' be two Riemannian metric fields on U . Let $\psi_0 \geq 1$ be a bound on the metric distortion. Suppose that U is included in a ball $B_g(p_0, r_0)$, with $p_0 \in U$ and $r_0 \in \mathbb{R}^+$, such that $\forall p \in B(p_0, r_0), \psi(g(p), g'(p)) \leq \psi_0$. Let \mathcal{P}_U be a point set in U . Suppose that \mathcal{P}_U is an (ε, μ) -net of U with respect to g . Then \mathcal{P}_U is an $(\varepsilon_{g'}, \mu_{g'})$ -net of U with respect to g' with $\varepsilon_{g'} = \psi_0 \varepsilon$ and $\mu_{g'} = \mu / \psi_0$.*

Proof. By Lemma B.1, we have that

$$\forall x, y \in U, \frac{1}{\psi_0} d_{g'}(x, y) \leq d_g(x, y) \leq \psi_0 d_{g'}(x, y).$$

Therefore

$$\forall x \in U, \exists p \in \mathcal{P}_U, d_g(x, p) \leq \varepsilon \Rightarrow \forall x \in U, \exists p \in \mathcal{P}_U, d_{g'}(x, p) \leq \psi_0 \varepsilon,$$

and

$$\forall p, q \in \mathcal{P}_U, d_g(p, q) \geq \mu \Rightarrow \forall p, q \in \mathcal{P}_U, d_{g'}(p, q) \geq \frac{\mu}{\psi_0}.$$

And, with $\varepsilon_{g'} = \psi_0 \varepsilon$ and $\mu_{g'} = \mu / \psi_0$, \mathcal{P} is an $(\varepsilon_{g'}, \mu_{g'})$ -net of U . □

Remark E.3. *We assumed that $\mu = \lambda \varepsilon$. By Lemma E.2, we have $\varepsilon' = \psi_0 \varepsilon$ and $\mu' = \frac{\mu}{\psi_0}$. Therefore*

$$\mu' = \frac{\lambda \varepsilon}{\psi_0} = \frac{\lambda}{\psi_0^2} \varepsilon'.$$

E.2.2 Stability of the power protection property

It is more complex to show that the assumption of power protection is preserved under metric perturbation. Previously, we have only considered two similar but arbitrary Riemannian metric fields g and g' on a neighborhood U . We now restrict ourselves to the case where g' is a uniform metric field. We shall always compare the metric field g in a neighborhood U with the uniform metric field $g' = g_0 = g(p_0)$ where $p_0 \in U$. Because g_0 and the Euclidean metric field differ only by a linear transformation, we simplify matters and assume that g_0 is the Euclidean metric field.

We now give conditions such that the point set is also protected with respect to g_0 . A few intermediary steps are needed to prove the main result:

- Given two sites, we prove that the bisectors of these two sites in the Voronoi diagrams built with respect to g and with respect to $g' = g_{\mathbb{E}}$ are close (Lemma E.5).
- We prove that the Voronoi cell of a point p_0 with respect to g , $V_g(p_0)$ can be encompassed by two scaled versions (one larger and one smaller) of the Euclidean Voronoi cell $V_{\mathbb{E}}(p_0)$ (Lemma E.12).
- We combine this encompassing with bounds on the dihedral angles of Euclidean Delaunay simplices given by Lemma D.6 to compute a stability region around Voronoi vertices where the same combinatorial Voronoi vertex of lives for both $V_g(p_0)$ and $V_{\mathbb{E}}(p_0)$ in 2D (Lemma E.13). We then extend it to any dimension by induction (Lemma E.14).

The main result appears in Lemma E.16 and gives the stability of power protection under metric perturbation.

We first define the scaled version of a Voronoi cell more rigorously.

Definition E.4 (Relaxed Voronoi cell). *Let $\omega \in \mathbb{R}$. The relaxed Voronoi cell of the site p_0 is*

$$V_g^{+\omega}(p_0) = \{x \in U \mid d_g(p_0, x)^2 \leq d_g(p_i, x)^2 + \omega \text{ for all } i \neq 0\}.$$

The following lemma expresses that two Voronoi cells computed in similar metric fields are close.

Lemma E.5. *Let $U \subset \Omega$ be open, and g and g' be two Riemannian metric fields on U . Let $\psi_0 \geq 1$ be a bound on the metric distortion. Suppose that U is included in a ball $B_g(p_0, r_0)$, with $p_0 \in U$ and $r_0 \in \mathbb{R}^+$, such that $\forall p \in B_g(p_0, r_0), \psi(g(p), g'(p)) \leq \psi_0$. Let $\mathcal{P}_U = \{p_i\}$ be a point set in U . Let $V_{p_0, g}$ denote a Voronoi cell with respect to the Riemannian metric field g .*

Suppose that the Voronoi cell $V_{g'}^{2\rho^2(\psi_0^2-1)}(p_0)$ lies in a ball of radius ρ with respect to the metric g' , which lies completely in U . Let $\omega_0 = 2\rho^2(\psi_0^2 - 1)$. Then $V_g(p_0)$ lies in $V_{g'}^{+\omega_0}(p_0)$ and contains $V_{g'}^{-\omega_0}(p_0)$.

Proof. Let $\text{BS}_g(p_0, p_i)$ be the bisector between p_0 and p_i with respect to g . Let $y \in \text{BS}_g(p_0, p_i) \cap B_{g'}(p_0, \rho)$, where $B_{g'}(p_0, \rho)$ denotes the ball centered at p_0 of radius ρ with respect to g' . Now $d_g(y, p_0) = d_g(y, p_i)$, and thus

$$\begin{aligned} |d_{g'}(y, p_0)^2 - d_{g'}(y, p_i)^2| &= |d_{g'}(y, p_0)^2 - d_g(y, p_0)^2 + d_g(y, p_i)^2 - d_{g'}(y, p_i)^2| \\ &\leq |d_{g'}(y, p_0)^2 - d_g(y, p_0)^2| + |d_{g'}(y, p_i)^2 - d_g(y, p_i)^2| \\ &\leq (\psi_0^2 - 1)(d_{g'}(y, p_0)^2 + d_{g'}(y, p_i)^2) \\ &\leq 2\rho^2(\psi_0^2 - 1). \end{aligned}$$

Thus $d_{g'}(y, p_0)^2 \leq d_{g'}(y, p_i)^2 + \omega$ and $d_{g'}(y, p_0)^2 \geq d_{g'}(y, p_i)^2 - \omega$ with $\omega = 2\rho^2(\psi_0^2 - 1)$, which gives us the expected result. \square

Remark E.6. Lemma E.5 does not require g' to be a uniform metric field.

We clarify in the next lemma that the bisectors of a Voronoi diagram with respect to a uniform are affine hyperplanes.

Lemma E.7. Let $U \subset \Omega$ be open, and g and g' be two Riemannian metric fields on U . Let $\psi_0 \geq 1$ be a bound on the metric distortion. Suppose that U is included in a ball $B_g(p_0, r_0)$, with $p_0 \in U$ and $r_0 \in \mathbb{R}^+$, such that $\forall p \in B(p_0, r_0), \psi(g(p), g'(p)) \leq \psi_0$. Let $\mathcal{P}_U = \{p_i\}$ be a point set in U . Let g' be a uniform metric field. We refer to g' as g_0 to emphasis its constancy. Let $p_0 \in \mathcal{P}_U$. The bisectors of $V_{g_0}^{\pm\omega_0}(p_0)$ are hyperplanes.

Proof. The bisectors of $V_{g_0}^{\pm\omega_0}(p_0)$ are given by

$$\text{BS}_{g_0}^{\pm\omega_0}(p_0, p_i) = \{x \in \Omega \mid d_{g_0}(p_0, x)^2 = d_{g_0}(p_i, x)^2 \pm \omega_0\}.$$

For $x \in \text{BS}_{g_0}^{\pm\omega_0}(p_0, p_i)$, we have by definition that

$$\begin{aligned} \|x - p_0\|_{g_0}^2 &= \|x - p_i\|_{g_0}^2 \pm \omega_0 \\ \iff (x - p_0)^t g_0 (x - p_0) &= (x - p_i)^t g_0 (x - p_i) \pm \omega_0 \\ \iff 2x^t g_0 (p_i - p_0) &= p_i^t g_0 p_i - p_0^t g_0 p_0 \pm \omega_0. \end{aligned}$$

which is the equation of an hyperplane since g_0 is uniform. \square

The cells $V_{g_0}^{\pm\omega_0}(p_0)$ are unfortunately impractical to manipulate as we do not have an explicit distance between the boundaries $\partial V_{g_0}(p_0)$ and $\partial V_{g_0}^{\pm\omega_0}(p_0)$. However that distance can be bounded; this is the purpose of the following lemma.

Lemma E.8. Let $U \subset \Omega$ be open, and g and g' be two Riemannian metric fields on U . Let $\psi_0 \geq 1$ be a bound on the metric distortion. Suppose that U is included in a ball $B_g(p_0, r_0)$, with $p_0 \in U$ and $r_0 \in \mathbb{R}^+$, such that $\forall p \in B(p_0, r_0), \psi(g(p), g'(p)) \leq \psi_0$. Let $\mathcal{P}_U = \{p_i\}$ be a point set in U . Assume furthermore that g_0 is the Euclidean metric field. Let $p_0 \in \mathcal{P}_U$. We have

$$d_{g_0}(\partial V_{g_0}(p_0), \partial V_{g_0}^{\pm\omega_0}(p_0)) := \min_{\substack{x \in \partial V_{g_0}(p_0) \\ y \in \partial V_{g_0}^{\pm\omega_0}(p_0)}} d_{g_0}(x, y) \leq \frac{\rho_0^2(\psi_0^2 - 1)}{\mu_0},$$

with ρ_0 defined as ρ is in Lemma E.5.

Proof. Let m_{ω_0} be the intersection of the segment $[p_0, p_i]$ and the bisector $\text{BS}_{g_0}^{-\omega_0}(p_0, p_i)$, for $i \neq 0$. Let m be the intersection of the segment $[p_0, p_i]$ and $\partial V_{g_0}(p_0)$, for $i \neq 0$. We have

$$\begin{cases} m_{\omega_0} \in \text{BS}_{g_0}^{-\omega_0}(p_0, p_i) & \iff 2m_{\omega_0}^T(p_i - p_0) = p_i^T p_i - p_0^T p_0 - \omega_0 \\ m \in \partial V_{g_0}(p_0) & \iff 2m^T(p_i - p_0) = p_i^T p_i - p_0^T p_0 \end{cases}$$

Therefore

$$2(m - m_{\omega_0})^T(p_i - p_0) = \omega_0.$$

Since $(m - m_{\omega_0})$ and $(p_i - p_0)$ are linearly dependent,

$$\omega_0 = 2|m - m_{\omega_0}||p_i - p_0|.$$

\mathcal{P} is μ_0 -separated, which implies that

$$\omega_0 \geq 2|m - m_{\omega_0}|\mu_0,$$

and

$$|m - m_{\omega_0}| \leq \frac{\omega_0}{2\mu_0} = \frac{\rho_0^2(\psi_0^2 - 1)}{\mu_0}.$$

□

Definition E.9. In the following, we use $\rho = 2\varepsilon_0$, and therefore $\omega_0 = 8\varepsilon_0^2(\psi_0^2 - 1)$. We show that this choice is reasonable in Lemma E.15. Additionally, we define

$$\eta_0 = \frac{\rho_0^2(\psi_0^2 - 1)}{\mu_0}.$$

We are now ready to encompass the Riemannian Voronoi cell of p_0 with respect to an arbitrary metric field g with two scaled versions of the Euclidean Voronoi cell of p_0 . The notions of *dilated* and *eroded* Voronoi cells will serve the purpose of defining precisely these scaled cells.

Definition E.10 (Eroded Voronoi cell). Let $\omega \in \mathbb{R}$. The eroded Voronoi cell of p_0 is the morphological erosion of $V_g(p_0)$ by a ball of radius ω :

$$EV_g^{-\omega}(p_0) = \{x \in V_g(p_0) \mid d_g(x, \partial V_g(p_0)) > \omega\}.$$

Definition E.11 (Dilated Voronoi cell). Let $\omega \in \mathbb{R}$. The dilated Voronoi cell of p_0 is:

$$DV_g^{+\omega}(p_0) = \bigcap_{i \neq 0} H^\omega(p_0, p_i),$$

where $H^\omega(p_0, p_i)$ is the half-space containing p_0 and delimited by the bisector $\mathcal{BS}(p_0, p_i)$ translated away from p_0 by ω_0 (see Figure 11).

The second important step on our path towards the stability of power protection is the encompassing of the Riemannian Voronoi cell, and is detailed below.

Lemma E.12. Let $U \subset \Omega$ be open, and g and g' be two Riemannian metric fields on U . Let $\psi_0 \geq 1$ be a bound on the metric distortion. Suppose that U is included in a ball $B_g(p_0, r_0)$, with $p_0 \in U$ and $r_0 \in \mathbb{R}^+$, such that $\forall p \in B(p_0, r_0), \psi(g(p), g'(p)) \leq \psi_0$. Let $\mathcal{P}_U = \{p_i\}$ be a point set in U . We have

$$EV_{g_0}^{-\eta_0}(p_0) \subseteq V_{g_0}^{-\omega_0}(p_0) \subseteq V_{g_0}(p_0) \subseteq V_{g_0}^{+\omega_0}(p_0) \subseteq DV_{g_0}^{+\eta_0}(p_0).$$

These inclusions are illustrated in Figure 11.

Proof. Using the notations and the result of Lemma E.8, we have

$$|m - m_{\omega_0}| \leq \eta_0.$$

Since the bisectors $\mathcal{BS}_{g'}^{\omega_0}(p_0, p_i)$ are hyperplanes, the result follows. □

In Figures 11 and 12, $DV_{g_0}^{+\eta_0}$ and $EV_{g_0}^{-\eta_0}$ are shown in green and V_{g_0} in yellow.

The next step is to prove that we have stability of the Voronoi vertices of $V_g(p_0)$, meaning that a small perturbation of the metric only creates a small displacement of any Voronoi vertex of the Voronoi cell of p_0 . The following lemma shows that Voronoi vertices are close if the distortion between the metric fields g and g_0 is small. The approach is to use Lemmas E.5 and E.12: we know that each $(n-1)$ -face of the Riemannian Voronoi cell $V_g(p_0)$ and shared by a Voronoi cell $V_g(q)$ is enclosed within the bisectors of $DV_{g_0}(p_0)$ and $EV_{g_0}(p_0)$ (which are translations of the




Figure 11: The different encompassing cells around p_0 . The Riemannian Voronoi diagrams with respect to g and g_0 are traced in dark cyan and black respectively. The Voronoi cell $V_{g_0}(p_0)$ is colored in yellow. The cells $DV_{g_0}^{+\eta_0}$ and $EV_{g_0}^{-\eta_0}$ are colored in green, and the cells $V_{g_0}^{\pm\omega_0}(p_0)$ are colored in purple.

bisectors of $V_{g_0}(p_0)$) for the sites p_0 and q . These two bisectors create a “band” that contains the bisector $BS_g(p_0, q)$. Given a Voronoi vertex c in $V_g(p_0)$, c can be obtained as the intersection of $n + 1$ Voronoi cells, but also as the intersection of n Voronoi $(n - 1)$ -faces of $V_g(p_0)$. The intersection of the bands associated to those n $(n - 1)$ -faces is a parallelotopic-shaped region which by definition contains the same (combinatorially speaking) Voronoi vertex, but for $V_{g_0}(p_0)$. Lemmas E.13 and E.14 express this reasoning, which is illustrated in Figure 12 for 2D and 13 for any dimension.

Lemma E.13. *We consider here $\Omega = \mathbb{R}^2$. Let $U \subset \Omega$ be open, and g and g' be two Riemannian metric fields on U . Let $\psi_0 \geq 1$ be a bound on the metric distortion. Suppose that U is included in a ball $B_g(p_0, r_0)$, with $p_0 \in U$ and $r_0 \in \mathbb{R}^+$, such that $\forall p \in B(p_0, r_0), \psi(g(p), g'(p)) \leq \psi_0$. Let $\mathcal{P}_U = \{p_i\}$ be a point set in U . Assume furthermore that \mathcal{P}_U is a δ_0 -power protected (ε_0, μ_0) -net with respect to g_0 (the Euclidean metric). Let $p_0 \in \mathcal{P}_U$. Let c and c_0 be the same Voronoi vertex in respectively $V_g(p_0)$ and $V_{g_0}(p_0)$. Then $d_{g_0}(c, c_0) \leq \chi_2$ with*

$$\chi_2 = \frac{2\eta_0}{\sqrt{1 + \frac{\mu_0}{2\varepsilon_0}} - \sqrt{1 - \frac{\mu_0}{2\varepsilon_0}}} = \frac{2\eta_0}{\sqrt{1 + \frac{\lambda}{2\psi_0^2}} - \sqrt{1 - \frac{\lambda}{2\psi_0^2}}}.$$

where λ is given by $\mu_0 = \frac{\lambda}{\psi_0^2}\varepsilon_0$ (see Lemma E.3).

Proof. We use Lemma E.12. $V_g(p_0)$ lies in $DV_{g_0}^{+\eta_0}$ and contains $EV_{g_0}^{-\eta_0}$. The circumcenters c and c_0 lie in a parallelogrammatic region centered on c_0 , itself included in the ball centered on c_0 and with radius χ . The radius χ is given by half the length of the longest diagonal of the parallelogram (see Figure 12). By Lemma E.2, \mathcal{P} is an (ε_0, μ_0) -net with respect to g_0 . Let θ be the angle of the Voronoi corner of $V_{g_0}(p_0)$ at c_0 . By Lemma D.1, that angle is bounded:

$$\theta_m = 2 \arcsin\left(\frac{\mu_0}{2\varepsilon_0}\right) \leq \theta \leq \pi - \arcsin\left(\frac{\mu_0}{2\varepsilon_0}\right) = \theta_M.$$

Since $\pi - \theta_M < \theta_m$, χ is maximal when $\theta > \pi/2$. We thus assume $\theta > \pi/2$, and compute a bound on χ as follows:

$$\begin{aligned} \sin\left(\frac{\pi - \theta}{2}\right) &= \frac{\eta_0}{\chi} \implies \chi = \frac{\eta_0}{\sin\left(\frac{\pi - \theta}{2}\right)} \\ &\leq \frac{\eta_0}{\sin\left(\frac{1}{2} \arcsin\left(\frac{\mu_0}{2\varepsilon_0}\right)\right)} \\ &\leq \frac{2\eta_0}{\sqrt{1 + \frac{\mu_0}{2\varepsilon_0}} - \sqrt{1 - \frac{\mu_0}{2\varepsilon_0}}} =: \chi_2, \end{aligned}$$

using $\sin(\frac{1}{2} \arcsin(\alpha)) = \frac{1}{2}(\sqrt{1 + \alpha} - \sqrt{1 - \alpha})$.




Figure 12: Black lines trace Vor_{g_0} and cyan lines trace Vor_g . The cell $V_{g_0}(p_0)$ is colored in yellow and the cell $V_g(p_0)$ is dashed. The green regions correspond to $DV_{g_0}^{+\eta_0}(p_0)$ and $EV_{g_0}^{-\eta_0}(p_0)$. On the left, the configuration where $\theta < \pi/2$; on the right, $\theta > \pi/2$.

□

The result obtained in Lemma E.13 can be extended to any dimension using induction and the stability of the Voronoi vertices of facets.

Lemma E.14. *We consider here $\Omega = \mathbb{R}^n$. Let $U \subset \Omega$ be open, and g and g' be two Riemannian metric fields on U . Let $\psi_0 \geq 1$ be a bound on the metric distortion. Suppose that U is included in a ball $B_g(p_0, r_0)$, with $p_0 \in U$ and $r_0 \in \mathbb{R}^+$, such that $\forall p \in B(p_0, r_0), \psi(g(p), g'(p)) \leq \psi_0$. Let $\mathcal{P}_U = \{p_i\}$ be a δ_0 -power protected (ε_0, μ_0) -net with respect to g_0 (the Euclidean metric). Let $p_0 \in \mathcal{P}_U$. Let c and c_0 be the same Voronoi vertex in respectively $V_g(p_0)$ and $V_{g_0}(p_0)$. Then $d_{g_0}(c, c_0) \leq \chi$ with*

$$\chi = \frac{\chi_2}{\sin^{n-2}\left(\frac{\varphi_0}{2}\right)}.$$

where χ_2 is defined as in Lemma E.13, and φ_0 is the maximal dihedral angle between two faces of a simplex.

Proof. We know from Lemma E.5 that $V_g(p_0)$ lies in $DV_{g_0}^{+\eta_0}$ and contains $EV_{g_0}^{-\eta_0}$. The circumcenters c and c_0 lie in a parallelotopic region centered on c_0 defined by the intersection of n Euclidean thickened Voronoi faces. This parallelotope and its circumscribing sphere are difficult to compute. However, it can be seen as the intersection of two parallelotopic tubes defined by the intersection of $n - 1$ Euclidean thickened Voronoi faces. From another point of view, this is

the computation of the intersection of the thickened duals of two facets τ_1 and τ_2 incident to p_0 of the simplex $\sigma \in Del(\mathcal{P})$, dual of c (and c_0), see Figure 13 (left).




Figure 13: Left, a simplex, the duals of two faces and their respective thickened duals (cylinders); the orange thickened dual is orthogonal to the purple face (and inversely). Right, the intersection of the two tubes, seen from above, illustrates the proof of Lemma E.14.

The stability radius χ is computed incrementally by increasing the dimension and proving stability of the circumcenters of the faces of the simplices. We prove the formula by induction.

The radius of each tube is given by the stability of the radius of the circumcenter in the lower dimension of the facet. The base case, $\mathbb{R}^n = \mathbb{R}^3$, is solved by Lemma E.13, and gives

$$\chi_2 = \frac{2\eta_0}{\sqrt{1 + \frac{\lambda}{2\psi_0^2}} - \sqrt{1 - \frac{\lambda}{2\psi_0^2}}}.$$

We now consider two facets τ_1 and τ_2 that are incident to p_0 . Denote \mathcal{D}_1 and \mathcal{D}_2 their respective duals. By Lemma E.5, \mathcal{D}_1 and \mathcal{D}_2 lie in two cylinders C_1 and C_2 of radius χ_2 . C_1 and C_2 are also orthogonal to τ_1 and τ_2 and c and c_0 lie in $C_1 \cap C_2$. The angle φ between C_1 and C_2 is exactly the dihedral angle between τ_1 and τ_2 . By Lemma D.6, we have

$$\arcsin(s_0) \leq \varphi \leq \pi - \arcsin(s_0) \text{ with } s_0 = \frac{1}{2} \left[\frac{\iota^2}{4\psi_0^2} - \frac{1}{2} \left(\psi_0^2 - \frac{1}{\psi_0^2} \right) \right]$$

Let $\varphi_0 = \arcsin(s_0) = \pi - \arcsin(s_0)$. We encompass the intersection of the cylinders, difficult to compute, with a sphere whose radius can be computed as follows (see Figure 13, right):

$$\chi_3 = \frac{\chi_2}{\cos(\alpha)} \text{ with } \alpha = \frac{\pi}{2} - \frac{\varphi_0}{2}.$$

Thus,

$$\chi_3 = \frac{\chi_2}{\sin\left(\frac{\varphi_0}{2}\right)}.$$

Recursively,

$$\chi = \frac{\chi_2}{\sin^{n-2}\left(\frac{\varphi_0}{2}\right)}.$$

□

We have assumed in different lemmas that we could pick values of ρ_0 or ω_0 that fit our need. The following lemma shows that these assumptions were reasonable.

Lemma E.15. *Lemmas E.12 and E.14 allow us to characterize the parameter ρ_0 more precisely. Indeed, an assumption of Lemma E.5 was that $V_{g_0}^{\omega_0}(p_0)$ is included in a ball $B_{g_0}(p_0, \rho_0)$. If the sampling of \mathcal{P} is sufficiently dense, such an assumption is reasonable.*

Proof. By definition, the Voronoi cell $V_{g_0}^{\omega_0}(p_0)$ is included in the dilated cell $DV_{g_0}^{+\eta_0}(p_0)$. Since the point set is an ε -sample, we have $d_{g_0}(p_0, x) \leq \varepsilon_0$ for $x \in V_{g_0}(p_0)$. By Lemma E.14, we have for $x \in DV_{g_0}^{+\eta_0}(p_0)$

$$d_{g_0}(p_0, x) \leq \varepsilon_0 + \chi.$$

Recall from Lemma E.14 that

$$\chi = \frac{\chi_2}{\left[\sin\left(\frac{1}{2} \arcsin(s_0)\right)\right]^{n-2}}.$$

with $\eta_0 = \frac{\rho_0^2(\psi_0^2-1)}{\mu_0}$ and $s_0 = \frac{1}{2} \left[\frac{\varepsilon^2}{4\psi_0^2} - \frac{1}{2} \left(\psi_0^2 - \frac{1}{\psi_0^2} \right) \right]$. We require $V_{g_0}^{\omega_0}(p_0) \subset B_{g_0}(p_0, \rho_0)$, which is verified if $DV_{g_0}^{+\eta_0}(p_0) \subset B_{g_0}(p_0, \rho_0)$, that is if

$$\begin{aligned} & \varepsilon_0 + \frac{\chi_2}{\left[\sin\left(\frac{1}{2} \arcsin(s_0)\right)\right]^{n-2}} \leq \rho_0 \\ \iff & \frac{4\eta_0}{\left(\sqrt{1 + \frac{\lambda}{2\psi_0^2}} - \sqrt{1 - \frac{\lambda}{2\psi_0^2}}\right) \left(\sqrt{1 + s_0} - \sqrt{1 - s_0}\right)^{n-2}} \leq \rho_0 - \varepsilon_0 \\ \iff & \frac{4\rho_0^2(\psi_0^2 - 1)}{\mu_0 \left(\sqrt{1 + \frac{\lambda}{2\psi_0^2}} - \sqrt{1 - \frac{\lambda}{2\psi_0^2}}\right) \left(\sqrt{1 + s_0} - \sqrt{1 - s_0}\right)^{n-2}} \leq \rho_0 - \varepsilon_0 \\ \iff & \frac{4(\psi_0^2 - 1)}{\left(\sqrt{1 + \frac{\lambda}{2\psi_0^2}} - \sqrt{1 - \frac{\lambda}{2\psi_0^2}}\right) \left(\sqrt{1 + s_0} - \sqrt{1 - s_0}\right)^{n-2}} \leq \frac{\mu_0(\rho_0 - \varepsilon_0)}{\rho_0^2} \\ \iff & \frac{4\psi_0(\psi_0^2 - 1)}{\left(\sqrt{1 + \frac{\lambda}{2\psi_0^2}} - \sqrt{1 - \frac{\lambda}{2\psi_0^2}}\right) \left(\sqrt{1 + s_0} - \sqrt{1 - s_0}\right)^{n-2}} \leq \frac{\lambda\varepsilon(\rho_0 - \varepsilon_0)}{\rho_0^2}, \text{ by Remark E.3.} \end{aligned}$$

The parameter ρ_0 can be chosen arbitrarily as long as it is greater than ε_0 , and we have taken $\rho_0 = 2\varepsilon_0$ (see Definition E.9), which imposes

$$\frac{\psi_0^2(\psi_0^2 - 1)}{\left(\sqrt{1 + \frac{\lambda}{2\psi_0^2}} - \sqrt{1 - \frac{\lambda}{2\psi_0^2}}\right) \left(\sqrt{1 + s_0} - \sqrt{1 - s_0}\right)^{n-2}} \leq \frac{\lambda}{16}. \quad (6)$$

Recall that the parameter λ is fixed. By continuity of the metric field, $\lim_{\varepsilon \rightarrow 0} \psi_0 = 1$, therefore the left hand side goes to 0 and Inequality (6) is eventually satisfied as the sampling is made denser. \square

Finally, we can now show the main result: the power protection property is preserved when the metric field is perturbed.

Lemma E.16. *Let $U \subset \Omega$ be open, and g and g' be two Riemannian metric fields on U . Let $\psi_0 \geq 1$ be a bound on the metric distortion. Suppose that U is included in a ball $B_g(p_0, r_0)$, with $p_0 \in U$ and $r_0 \in \mathbb{R}^+$, such that $\forall p \in B(p_0, r_0), \psi(g(p), g'(p)) \leq \psi_0$. Assume that \mathcal{P}_U is a δ -power protected (ε, μ) -net in U with respect to g . If δ is well chosen, then \mathcal{P}_U is a δ_0 -power protected net with respect to g_0 , with*

$$\delta_0^2 = \left(\frac{1}{\psi_0^2} - \psi_0^2 \right) (\varepsilon + \chi)^2 - \frac{4\varepsilon\chi}{\psi_0^2} + \frac{\delta^2}{\psi_0^2}.$$

Proof. By Lemma E.2, we know that \mathcal{P}_U is (ε_0, μ_0) -net with respect to g_0 . Let $q \in \mathcal{P}_U$, with q not a vertex in the dual of c . Let c_0 be the combinatorial equivalent of c in $V_{g_0}(\mathcal{P})$. Since \mathcal{P} is a δ -power protected net with respect to g , we have $d_g(c, q) > \sqrt{r^2 + \delta^2}$, where $r = d_g(c, p)$. On the one hand, we have

$$\begin{aligned} d_{g_0}(c_0, q) &\geq d_{g_0}(q, c) - d_{g_0}(c, c_0) \\ &\geq \frac{1}{\psi_0} d_g(q, c) - \chi \\ &\geq \frac{1}{\psi_0} \sqrt{r^2 + \delta^2} - \chi. \end{aligned}$$

by Lemma E.13. On the other hand, for any $p \in \mathcal{P}_U$ such that p is a vertex of the dual of c , we have

$$\begin{aligned} r_0 = d_{g_0}(c_0, p) &\leq d_{g_0}(c, p) + d_{g_0}(c_0, c) \\ &\leq \psi_0 d_g(c, p) + \chi \\ &\leq \psi_0 r + \chi. \end{aligned}$$

Thus δ_0 -power protection of \mathcal{P}_U with respect to g_0 requires

$$\begin{aligned} \frac{1}{\psi_0} \sqrt{r^2 + \delta^2} - \chi &> \chi + \psi_0 r \\ \iff \sqrt{r^2 + \delta^2} &> \psi_0(2\chi + \psi_0 r). \end{aligned}$$

This is verified if

$$\begin{aligned} \delta^2 &> (\psi_0(2\chi + \psi_0 r))^2 - r^2 = 4\chi^2\psi_0^2 + 4\chi\psi_0^3r + \psi_0^4r^2 - r^2 \\ &= 4\chi^2\psi_0^2 + 4\chi\psi_0^3r + (\psi_0^4 - 1)r^2, \end{aligned}$$

for all $r \in [\mu/2, \varepsilon]$. This gives us

$$\delta^2 > 4\chi^2\psi_0^2 + 4\chi\psi_0^3\varepsilon + (\psi_0^4 - 1)\varepsilon^2. \quad (7)$$

This condition on δ is only reasonable if the right hand side is not too large. Indeed, since \mathcal{P} is an ε -sample, we must have $d_g(c, q) < 2\varepsilon$. However, we have that $d_g(c, q)^2 > d_g(c, p)^2 + \delta^2$ by δ -power protection of \mathcal{P} with respect to g . Because $d_g(c, p) < \varepsilon$, it suffices that $\delta < \varepsilon$. We will now show that this is reasonable by examining the limit of the right hand side of Inequality (7).

We note, see Lemma E.14, that

$$\chi = \frac{\chi_2}{\sin^{n-2}(\frac{\varphi}{2})} = \frac{4\eta}{\left(\sqrt{1 + \frac{\lambda}{2\psi_0^2}} - \sqrt{1 - \frac{\lambda}{2\psi_0^2}}\right) \left(\sqrt{1 + s_0} - \sqrt{1 - s_0}\right)^{n-2}},$$

where $\varepsilon_0 = \psi_0\varepsilon$ and $\mu_0 = \mu/\psi_0 = \lambda\varepsilon/\psi_0$ (see Remark E.3). So that

$$\begin{aligned} 4\chi^2\psi_0^2 + 4\chi\psi_0^3\varepsilon + (\psi_0^4 - 1)\varepsilon^2 &= 4 \left(\frac{16\varepsilon\psi_0^4(\psi_0^2 - 1)}{\lambda \left(\sqrt{1 + \frac{\lambda}{2\psi_0^2}} - \sqrt{1 - \frac{\lambda}{2\psi_0^2}}\right) \left(\sqrt{1 + s_0} - \sqrt{1 - s_0}\right)^{n-2}} \right)^2 \\ &\quad + 4 \frac{16\varepsilon\psi_0^3(\psi_0^2 - 1)}{\lambda \left(\sqrt{1 + \frac{\lambda}{2\psi_0^2}} - \sqrt{1 - \frac{\lambda}{2\psi_0^2}}\right) \left(\sqrt{1 + s_0} - \sqrt{1 - s_0}\right)^{n-2}} \psi_0^3 \\ &\quad + (\psi_0^2 - 1)(\psi_0^2 + 1)\varepsilon^2. \end{aligned} \quad (8)$$

This means that the right hand side of (7) is of the form $f(\psi_0)(\psi_0^2 - 1)\varepsilon^2$, where $f(\psi_0)$ is a function that tends to a constant as ψ_0 goes to 1:

$$\begin{aligned}
 f(\psi_0) &= 4 \left(\frac{16\psi_0^4}{\lambda \left(\sqrt{1 + \frac{\lambda}{2\psi_0^2}} - \sqrt{1 - \frac{\lambda}{2\psi_0^2}} \right) \left(\sqrt{1 + s_0} - \sqrt{1 - s_0} \right)^{n-2}} \right)^2 \\
 &\quad + 4 \frac{16\psi_0^6}{\lambda \left(\sqrt{1 + \frac{\lambda}{2\psi_0^2}} - \sqrt{1 - \frac{\lambda}{2\psi_0^2}} \right) \left(\sqrt{1 + s_0} - \sqrt{1 - s_0} \right)^{n-2}} \\
 &\quad + (\psi_0^2 + 1) \\
 &\xrightarrow{\psi_0 \rightarrow 1} 4 \left(\frac{16}{\lambda \left(\sqrt{1 + \frac{\lambda}{2}} - \sqrt{1 - \frac{\lambda}{2}} \right) \left(\sqrt{1 + \frac{\varepsilon^2}{4}} - \sqrt{1 - \frac{\varepsilon^2}{4}} \right)^{n-2}} \right)^2 \\
 &\quad + 4 \frac{16}{\lambda \left(\sqrt{1 + \frac{\lambda}{2}} - \sqrt{1 - \frac{\lambda}{2}} \right) \left(\sqrt{1 + \frac{\varepsilon^2}{4}} - \sqrt{1 - \frac{\varepsilon^2}{4}} \right)^{n-2}} \\
 &\quad + 2
 \end{aligned}$$

So the bound given in Equality (7) may be easily satisfied if the metric distortion is sufficiently small.

We now provide an explicit value for δ_0 in terms of δ . Let $\xi = d_g(c, q)$ and $\xi_0 = d_{g_0}(c_0, q_0)$. We have the following bounds on r_0 and ξ_0 :

$$\begin{aligned}
 \frac{1}{\psi_0}(r - \xi) &\leq r_0 \leq \psi_0(r + \xi) \\
 \frac{1}{\psi_0} \sqrt{(r - \chi)^2 + \delta^2} &\leq \xi_0 \leq \psi_0 \sqrt{(r + \chi)^2 + \delta^2}.
 \end{aligned}$$

If we had $\tilde{\delta}$ -power protection, we would have

$$\begin{aligned}
 r_0^2 + \tilde{\delta}^2 \leq \xi_0^2 &\iff \tilde{\delta}^2 \leq \xi_0^2 - r_0^2 \\
 &\iff \tilde{\delta}^2 \leq \frac{1}{\psi_0^2} ((r - \chi)^2 + \delta^2) - \psi_0^2(r + \chi)^2 \\
 &\iff \tilde{\delta}^2 \leq \frac{1}{\psi_0^2}(r + \chi)^2 - \frac{4r\chi}{\psi_0^2} + \frac{\delta^2}{\psi_0^2} - \psi_0^2(r + \chi)^2 \\
 &\iff \tilde{\delta}^2 \leq \left(\frac{1}{\psi_0^2} - \psi_0^2 \right) (\varepsilon + \chi)^2 - \frac{4\varepsilon\chi}{\psi_0^2} + \frac{\delta^2}{\psi_0^2}.
 \end{aligned}$$

Therefore we can take $\delta_0^2 = \frac{\delta^2}{\psi_0^2} + \left(\frac{1}{\psi_0^2} - \psi_0^2 \right) (\varepsilon + \chi)^2 - \frac{4\varepsilon\chi}{\psi_0^2}$. Note that with this definition, δ_0 goes to δ as ψ_0 goes to 1, which proves that our value of δ_0 is legitimate. \square

F Embeddability of the straight Delaunay triangulation (Proofs of Section 3.2)

We first prove Lemma 3.3, recalled below, which bounds the distance between the same point on a the Karcher and the straight simplex.

Lemma. Let \mathcal{P} be an ε -sample with respect to g on Ω . Let $\{p_i\}$ be a set of $n+1$ vertices in \mathcal{M} such that $N = \cap_{p_i \in \mathcal{P}} V(p_i) \neq \emptyset$. Let $\bar{\sigma}$ and $\tilde{\sigma}$ be the straight and Karcher simplices that realize N . Let \tilde{x} be a point on the Karcher simplex $\tilde{\sigma}$ determined by the barycentric coordinates $\{\lambda_i\}$ (see Equation 2). Let x_e be the point uniquely determined by $\{\lambda_i\}$ as $x_e = \sum_i \lambda_i p_i$. If the geodesic distance d_g is close to $d_{\mathbb{E}}$ – that is the distortion $\psi(g, g_{\mathbb{E}})$ is bounded by ψ_0 – then $|\tilde{x} - x_e| \leq \sqrt{2 \cdot 4^3(\psi_0 - 1)\varepsilon^2}$.

Proof. The key observation is that given a convex function f and a function f' that is close, that is $f - f' < \alpha$ with α small, then the minimum value of f' is at most of $\min f + \alpha$. If we observe that at any point x where $f(x) > \min f + 2\alpha$, we also have $f'(x) > \min f + \alpha$ so x is not a minimum of f' , we see that the minima of f and f' can not be far apart. In particular, we have that if $x_{f', \min}$ is the point where f' attains its minimum, then $f'(x_{f', \min}) \leq \min f + \alpha$. The precise argument goes as follows.

We again assume that (possibly after a linear transformation) the metric is close to the Euclidean one, that is:

$$d_g(x, y) = |x - y| + \delta d_g(x, y),$$

with $|\delta d_g(x, y)| \leq (\psi_0 - 1)|x - y|$. If we assume that $|x - y| \leq 4\varepsilon$ and $\psi_0 \leq 2$, it follows that

$$d_g(x, y)^2 = |x - y|^2 + \delta d_g^2(x, y),$$

with

$$\delta d_g^2(x, y) \leq 4^3(\psi_0 - 1)\varepsilon^2.$$

Recall that \tilde{x} is the point where the functional

$$\mathcal{E}_\lambda(x) = \sum_i \lambda_i d_g(x, p_i)^2$$

attains its minimum.

Using the bounds above, we find that

$$\sum_i \lambda_i d_g(x, p_i)^2 = \sum_i \lambda_i |x - p_i|^2 + \sum_i \lambda_i \delta d_g^2(x, p_i),$$

where $\sum_i \lambda_i \delta d_g^2(x, p_i) \leq 4^3(\psi_0 - 1)\varepsilon^2$. We also see that

$$\left| \sum_i \lambda_i d_g(\tilde{x}, p_i)^2 - \sum_i \lambda_i |\tilde{x} - p_i|^2 \right| \leq 4^3(\psi_0 - 1)\varepsilon^2.$$

Taking f' to be $\sum_i \lambda_i d_g(\tilde{x}, p_i)^2$ in the explanation above, we find that

$$\left| \sum_i \lambda_i d_g(\tilde{x}, p_i)^2 - \sum_i \lambda_i \left| \sum_j \lambda_j v_j - p_i \right|^2 \right| \leq 4^3(\psi_0 - 1)\varepsilon^2,$$

because the Euclidean barycenter $x_e = \sum_i \lambda_i p_i$ is where the function $\sum_i \lambda_i |\tilde{x} - p_i|^2$ attains its minimum. Combining these results yields

$$\left| \sum_i \lambda_i |\tilde{x} - p_i|^2 - \sum_i \lambda_i \left| \sum_j \lambda_j v_j - p_i \right|^2 \right| \leq 2 \cdot 4^3(\psi_0 - 1)\varepsilon^2. \quad (9)$$

This bounds the distance between \tilde{x} and x_e . An explicit bound can be found by observing that

$$\begin{aligned}
\sum_i \lambda_i |x - p_i|^2 &= \sum_i \lambda_i \left((x^1 - p_i^1)^2 + \dots + (x^n - p_i^n)^2 \right) \\
&= \sum_i \lambda_i (x^1 - p_i^1)^2 + \sum_i \lambda_i (x^2 - p_i^2)^2 + \dots + \sum_i \lambda_i (x^n - p_i^n)^2 \\
&= \left(x^1 - \sum_i \lambda_i p_i^1 \right)^2 - \left(\sum_i \lambda_i p_i^1 \right)^2 + \sum_i \lambda_i (p_i^1)^2 + \dots \\
&\quad + \left(x^n - \sum_i \lambda_i p_i^n \right)^2 - \left(\sum_i \lambda_i p_i^n \right)^2 + \sum_i \lambda_i (p_i^n)^2 \\
&= \left| x - \sum_i \lambda_i p_i \right|^2 + \sum_j \left[- \left(\sum_i \lambda_i p_i^j \right)^2 + \sum_i \lambda_i (p_i^j)^2 \right]. \tag{10}
\end{aligned}$$

Then, applying Equation 10 for both $x = \tilde{x}$ and $x = x_e = \sum_j \lambda_j p_j$ in Equation (9), we obtain:

$$\begin{aligned}
&\left| \left| \tilde{x} - \sum_i \lambda_i p_i \right|^2 + \sum_j \left[- \left(\sum_i \lambda_i p_i^j \right)^2 + \sum_i \lambda_i (p_i^j)^2 \right] \right. \\
&\quad \left. - \left(\left| \sum_j \lambda_j p_j - \sum_i \lambda_i p_i \right|^2 + \sum_j \left[- \left(\sum_i \lambda_i p_i^j \right)^2 + \sum_i \lambda_i (p_i^j)^2 \right] \right) \right| \\
&= \left| \tilde{x} - \sum_i \lambda_i p_i \right|^2 \leq 2 \cdot 4^3 (\psi_0 - 1) \varepsilon^2.
\end{aligned}$$

which yields a distance bound of $\sqrt{2 \cdot 4^3 (\psi_0 - 1) \varepsilon^2}$. \square

Although we have formulated this metric distortion result for simplices, the same proof extends almost verbatim to continuous distributions. By this we mean that the barycenter with respect to a metric g of a continuous distribution is close to the barycenter with respect to the Euclidean metric, if g is close to the Euclidean metric. Furthermore, note that the proof does not depend on the weights being positive.

We now prove Theorem 3.4, recalled below.

Theorem. *Let \mathcal{P} be a δ -power protected (ε, μ) -net with respect to g on Ω . Let $\{p_i\}$ be a set of $n + 1$ vertices in Ω such that $\cap_{p_i \in \mathcal{P}} V(p_i) \neq \emptyset$. Let $\bar{\sigma}$ and $\tilde{\sigma}$ be the straight and Karcher simplices with vertices $\{p_i\}$. Let $\tilde{\tau}$ be a facet of $\tilde{\sigma}$, opposite of the vertex p_i . If for all $\tilde{x} \in \tilde{\tau}$, we have $|\tilde{x} - x_e|$ smaller than the lower bound on $D(p_i, \sigma)$, where x_e is the corresponding point on $\bar{\sigma}$ (as defined in Lemma 3.3), then there is no inversion created when $\tilde{\sigma}$ is straightened onto $\bar{\sigma}$. Furthermore, if this condition is fulfilled for all $\tilde{\sigma} \in \widetilde{\text{Del}}_g(\mathcal{P})$, then $\overline{\text{Del}}_g(\mathcal{P})$ is embedded.*

Proof. The lower bound on $D(p, \sigma)$ given in Appendix D is proportional to ε . The proximity (upper) bound from Lemma 3.3 is proportional to $\sqrt{(\psi_0 - 1)\varepsilon}$, therefore going to 0 much faster. The embeddability is thus satisfied once

$$\sqrt{2 \cdot 4^3 (\psi_0 - 1) \varepsilon^2} < \frac{\delta^2}{4\varepsilon} \iff \psi_0 < 1 + \frac{\iota^4}{32 \cdot 4^3}.$$

\square

G Deforming the triangulation \mathcal{T}_v

An extreme configuration can have a sphere in Figure 14 separating parts of \mathcal{T}_v from the Voronoi vertex v . In that case, we can chose different spheres to “push away” faces. Figure 14 shows the construction. We do not detail the computations.




Figure 14: Limit case of the deformation of \mathcal{T}_v into \mathcal{T}_v .

H Equality of the Riemannian Delaunay complexes in advanced geometric settings

In this section, we explain precisely how to obtain conditions on the maximal length of a canvas edge and the quality of the sample set such that the $\text{Del}_g^d(\mathcal{P})$ and the $\text{Del}_g(\mathcal{P})$ are equal.

H.1 Uniform metric field

We first investigate the setting of a subdomain of \mathbb{R}^n endowed with a uniform metric field.

If, for all $\sigma \in \text{Del}_g^d(\mathcal{P})$, there exists $\sigma_C \in \mathcal{C}$ such that σ_C witnesses σ , and $\forall \sigma_C \in \mathcal{C}$ the simplex witnessed by σ_C belongs to $\text{Del}_g^d(\mathcal{P})$, then we say that \mathcal{C} captures $\text{Del}_g^d(\mathcal{P})$. A Voronoi cell is said to be *captured* if all its Voronoi vertices are witnessed by the canvas.

By Lemma E.1, if a point set \mathcal{P} in \mathbb{R}^n is a δ -power protected (ε, μ) -net with respect to a uniform metric field g_0 then the point set $\mathcal{P}' = \{F_0 p_i, p_i \in \mathcal{P}\}$, where F_0 is the square root of g_0 , is also a δ -power protected (ε, μ) -net, but with respect to the Euclidean metric. We can deduce

an upper bound on the maximal length of an edge of \mathcal{C} for a uniform metric field using the results of Theorem 6.1 for the Euclidean setting. The main result is given by the following theorem:

Theorem H.1. *Let \mathcal{P} be a point set in Ω . Let g be a uniform Riemannian metric field on Ω ($\forall x \in \Omega, g(x) = g_0$). Let \mathcal{C} be the canvas and let $e_{\mathcal{C}}$ be the length of its longest edge. If*

$$e_{\mathcal{C}} < \left(\min_i \sqrt{\lambda_i} \right) \min \left\{ \frac{\mu}{3}, \frac{\delta^2}{32\varepsilon} \right\},$$

then $\text{Del}_g^d(\mathcal{P}) = \text{Del}_g(\mathcal{P})$.

Proof. Consider $\mathcal{R} = \text{Vor}_{g_0}^d(\mathcal{P})$ over \mathcal{C} . Let F_0 be a square root of G_0 . The matrix F_0 provides a stretching operator between the Euclidean and the metric spaces. Let $\mathcal{P}_0 = \{F_0 p, p \in \mathcal{P}\}$ be the transformed point set and \mathcal{C}_0 be a canvas (a dense triangulation) of the transformed space. Denote by \mathcal{R}_0 the discrete Riemannian Voronoi diagram of \mathcal{P}_0 with respect to $g_{\mathbb{E}}$ over \mathcal{C}_0 . Let $e_{\mathcal{C},0}$ be the upper bound on the canvas edge length of \mathcal{C}_0 provided by Theorem 6.1 such that \mathcal{D}_0 is captured by \mathcal{C}_0 .

Since \mathcal{P}_0 is a δ -power protected net with respect to $g_{\mathbb{E}}$, we can invoke Theorem 6.1 and we must thus have

$$e_{\mathcal{C},0} < \min \left\{ \frac{\mu}{3}, \frac{\delta^2}{32\varepsilon} \right\},$$

for the canvas \mathcal{C}_0 to capture $\text{Del}_{\mathbb{E}}(\mathcal{P}_0)$.

Let \mathcal{C}'_0 be the image of \mathcal{C} by F_0 . Note that \mathcal{C}'_0 and \mathcal{C}_0 are two different triangulations of the same space. If any edge of \mathcal{C}'_0 is smaller (with respect to the Euclidean metric) than $e_{\mathcal{C},0}$, then \mathcal{C}'_0 satisfies

$$e_{\mathcal{C}',0} < \min \left\{ \frac{\mu}{3}, \frac{\delta^2}{32\varepsilon} \right\}$$

and thus \mathcal{C}'_0 captures $\text{Del}_{\mathbb{E}}(\mathcal{P}_0)$.

Recall that given an eigenvector v_i of G_0 with corresponding eigenvalue λ_i , a unit length in the direction v_i in the metric space has length $1/\sqrt{\lambda_i}$ in the Euclidean space. Therefore, if the bound $e_{\mathcal{C}}$ of \mathcal{C} is smaller than $\alpha e_{\mathcal{C},0}$, with

$$\alpha = \frac{1}{\max_i \left(\frac{1}{\sqrt{\lambda_i}} \right)} = \min_i \sqrt{\lambda_i},$$

then every edge of \mathcal{C}'_0 is smaller than $e_{\mathcal{C},0}$. This implies that \mathcal{C}'_0 captures $\text{Del}_{\mathbb{E}}(\mathcal{P}_0)$ and therefore that \mathcal{C} captures $\text{Del}_{g_0}(\mathcal{P})$. \square

This settles the case of a uniform metric field.

H.2 Arbitrary metric field

We now consider an arbitrary metric field g over the domain \mathbb{R}^n . The key to proving the equality of the discrete Riemannian Delaunay complex and the Riemannian Delaunay complex in this setting is to locally approximate the arbitrary metric field with a uniform metric field, a configuration that we have dealt with in the previous section. We shall always compare the metric field g in a neighborhood U with the uniform metric field $g' = g(p_0)$ where $p_0 \in U$. Because g' and the Euclidean metric field differ by a linear transformation, we can simplify matters and assume that g' is the Euclidean metric field. The main argument of the proof will be once again that a power protected has stable and separated Voronoi vertices.

We recall the main result of this section, Theorem 5.1.

Theorem. *Let g be an arbitrary metric field on Ω . Assume that \mathcal{P} is a δ -power protected (ε, μ) -net in Ω with respect to g . Denote by \mathcal{C} the canvas, and $e_{\mathcal{C}}$ the length of its longest edge. If*

$$e_{\mathcal{C}} < \min_{p \in \mathcal{P}} e_{\mathcal{C},p},$$

where $e_{\mathcal{C},p}$ is given by Lemma H.2, and if ε is sufficiently small and δ is sufficiently large (both values will be detailed in the proof), then $\text{Del}_g^d(\mathcal{P}) = \text{Del}_g(\mathcal{P})$.

We prove Theorem 5.1 by computing for each point $p \in \mathcal{P}$ the maximal edge length of the canvas such that the Voronoi cell $V_g(p)$ is captured correctly. Conditions on ε and δ shall emerge from the intermediary results on the stability of power protected nets.

Lemma H.2. *Let $\psi_0 \geq 1$ be a bound on the metric distortion and g be a Riemannian metric field on U . Let U be an open neighborhood of $\Omega = \mathbb{R}^n$ that is included in a ball $B_g(p_0, r_0)$, with $p_0 \in U$ and $r_0 \in \mathbb{R}^+$ such that $\forall p \in B(p_0, r_0), \psi(g(p), g_{\mathbb{E}}(p)) \leq \psi_0$. Let \mathcal{P}_U be a point set in U and let $p_0 \in \mathcal{P}_U$.*

Suppose that \mathcal{P}_U is a δ -power protected (ε, μ) -net of with respect to g . Let $V_g(p_0)$ be the Voronoi cell of p_0 in $\text{Vor}_g^d(\mathcal{P}_U)$. If

$$e_{\mathcal{C},p_0} < \min_i \left(\sqrt{\lambda_i} \right) \min \left\{ \frac{\mu}{3}, \frac{\ell_0}{2} \right\},$$

with $\{\lambda_i\}$ the eigenvalues of g_0 and ℓ_0 that is made explicit in the proof, and if ε is sufficiently small and δ is sufficiently large (both values will also be detailed in the proof), then $\text{Del}_g^d(\mathcal{P}) = \text{Del}_g(\mathcal{P})$.

H.2.1 Approach

The many intermediary results needed to prove Theorem 5.1 are presented in Appendices C, D and E. We refer to them at appropriate times.

We use the fact that a Riemannian Voronoi cell $V_g(p_0)$ can be encompassed into two Euclidean Voronoi cells $DV_{\mathbb{E}}^{+\eta}(p_0)$ and $EV_{\mathbb{E}}^{-\eta}(p_0)$ that are scaled up and down versions of $V_{\mathbb{E}}(p_0)$ (Lemma E.12). Specifically, $EV_{\mathbb{E}}^{-\eta}(p_0)$ and $DV_{\mathbb{E}}^{+\eta}(p_0)$ are defined by

$$EV_{\mathbb{E}}^{-\omega}(p_0) = \{x \in V_{\mathbb{E}}(p_0) \mid d_{\mathbb{E}}(x, \partial V_{\mathbb{E}}(p_0)) > \omega\},$$

and

$$DV_{\mathbb{E}}^{+\omega}(p_0) = \bigcap_{i \neq 0} H^{\omega}(p_0, p_i),$$

where $H^{\omega}(p_0, p_i)$ is the half-space containing p_0 and delimited by the bisector $\text{BS}(p_0, p_i)$ translated away from p_0 by ω_0 . The constant η is the thickness of this encompassing and depends on the bound on the distortion ψ_0 in the neighborhood, and on the sampling and separation parameters ε and μ . We have that η goes to 0 as ψ_0 goes to 1.

The (local) stability of the power protected nets assumption is proved here again (Lemmas E.2 and E.16). From this observation, we can deduce that the Voronoi vertices of the Euclidean Voronoi cell $V_{\mathbb{E}}(p_0)$ are separated, and thus that the Voronoi vertices of $EV_{\mathbb{E}}(p_0)$ are separated. A bound on the maximal length of a canvas edge can then be computed such that $EV_{\mathbb{E}}(p_0)$ is captured and thus $V_g(p_0)$ is captured.

Difficulties The difficulty almost entirely comes from proving the stability of the assumption of power protection under metric perturbation in any dimension, that is proving that if we assume that \mathcal{P} is δ -power protected with respect to the arbitrary metric field g , then, in a small neighborhood around p_0 , the point set \mathcal{P} is δ_0 -power protected with respect to $g_0 = g(p)$. Assuming a power protected net does give us some bounds, but creates a tricky circular dependency as the coefficient δ appears in the dihedral angles (see Lemma D.2). We remedy this issue by proving that Euclidean dihedral angles are bounded assuming power protection with respect to the arbitrary metric field, with Lemmas D.5 and D.6.

H.2.2 Proof of Lemma H.2

Lemma E.5 gives us that $V_g(p_0)$ lies in $DV_{\mathbb{E}}^{+\eta}(p_0)$ and contains $EV_{\mathbb{E}}^{-\eta}(p_0)$. Since $V_g(p_0)$ contains $EV_{\mathbb{E}}^{-\eta}(p_0)$, if e_c is small enough such that $EV_{\mathbb{E}}^{-\eta}(p_0)$ is captured, then $V_g(p_0)$ is also captured. Proving that $EV_{\mathbb{E}}^{-\eta}(p_0)$ is captured is done similarly to the Euclidean setting. While we do not explicitly have the power protected net property for the relaxed Voronoi cells (and specifically, $EV_{\mathbb{E}}^{-\eta}(p_0)$), we can still extract the critical property that the Voronoi vertices are separated, as shown by the next lemma.

Lemma H.3. *Assume U , g , and ψ_0 as in Lemma H.2. Assume that the point set \mathcal{P}_U is a δ_0 -power protected (ε, μ) -net with respect to the Riemannian metric field g . Then the Voronoi vertices of $EV_{\mathbb{E}}^{-\eta}(p_0)$ are separated.*

Proof. By Lemmas E.2 and E.16, we have local stability of the power protection and net properties. Hence, \mathcal{P} is δ_0 -power protected (ε_0, μ_0) -net with respect to $g_{\mathbb{E}}$ in U .

Let $L_0 = \delta_0^2/4\varepsilon_0$ be the separation bound induced by the δ_0 -power protection property of \mathcal{P}_U (see Lemma C.4). Let l be the distance between any two adjacent Voronoi vertices of $EV_{\mathbb{E}}^{-\eta}(p_0)$. We know by Lemma E.14 that the parallelotopic region around a Voronoi vertex is included in a ball centered on the Voronoi vertex and of radius χ . The protection parameter ι is given by $\delta = \iota\varepsilon$. We have that

$$\begin{aligned} l &\geq L - 2\chi \\ &\geq \frac{\delta_0^2}{4\varepsilon_0} - 2 \frac{\chi^2}{\sin^{n-2}\left(\frac{\varphi}{2}\right)} \\ &\geq \frac{\delta_0^2}{4\varepsilon_0} - 2 \frac{4\eta}{\left(\sqrt{1 + \frac{\lambda}{2\psi_0^2}} - \sqrt{1 - \frac{\lambda}{2\psi_0^2}}\right) (\sqrt{1 + s_0} - \sqrt{1 - s_0})^{n-2}} =: \ell_0, \end{aligned} \quad (11)$$

where φ represents the dihedral angle and $s_0 = \frac{1}{2} \left[\frac{\iota^2}{4\psi_0^2} - \frac{1}{2} \left(\psi_0^2 - \frac{1}{\psi_0^2} \right) \right]$. For the stability regions not to intersect, we require l to be positive. This can be ensured by enforcing that the lower bound is positive:

$$\frac{\delta_0^2}{4\varepsilon_0} > \frac{8\eta}{\left(\sqrt{1 + \frac{\lambda}{2\psi_0^2}} - \sqrt{1 - \frac{\lambda}{2\psi_0^2}}\right) (\sqrt{1 + s_0} - \sqrt{1 - s_0})^{n-2}}$$

Recall that $\varepsilon_0 = \psi_0\varepsilon$, $\mu_0 = \lambda\varepsilon/\psi_0$ and $\rho_0 = 2\psi_0\varepsilon$. Using these notations, we see that $l > 0$ if

$$\begin{aligned} \delta_0^2 &> \frac{32\varepsilon_0\rho_0^2(\psi_0^2 - 1)}{\mu_0 \left(\sqrt{1 + \frac{\lambda}{2\psi_0^2}} - \sqrt{1 - \frac{\lambda}{2\psi_0^2}}\right) (\sqrt{1 + s_0} - \sqrt{1 - s_0})^{n-2}} \\ \delta_0^2 &> \frac{128\varepsilon^2\psi_0^4(\psi_0^2 - 1)}{\lambda \left(\sqrt{1 + \frac{\lambda}{2\psi_0^2}} - \sqrt{1 - \frac{\lambda}{2\psi_0^2}}\right) (\sqrt{1 + s_0} - \sqrt{1 - s_0})^{n-2}}. \end{aligned} \quad (12)$$

This condition is easy to satisfy when ψ_0 goes to 1 because the right hand side of Inequality (12) is proportional to $(\psi_0^2 - 1)\varepsilon^2$.

The intermediary results that we use already impose some conditions on δ and ε , and we thus would like to give the condition in Equation 12 in terms of δ , so that it may be compared with Inequality (7). In Lemma E.16, we have seen that

$$\delta_0^2 = \frac{\delta^2}{\psi_0^2} + \left(\frac{1}{\psi_0^2} - \psi_0^2 \right) (\varepsilon + \chi)^2 - \frac{4\varepsilon\chi}{\psi_0^2},$$

Thus

$$\frac{\delta^2}{\psi_0^2} + \left(\frac{1}{\psi_0^2} - \psi_0^2 \right) (\varepsilon + \chi)^2 - \frac{4\varepsilon\chi}{\psi_0^2} > \frac{128\varepsilon^2\psi_0^4(\psi_0^2 - 1)}{\lambda \left(\sqrt{1 + \frac{\lambda}{2\psi_0^2}} - \sqrt{1 - \frac{\lambda}{2\psi_0^2}} \right) (\sqrt{1 + s_0} - \sqrt{1 - s_0})^{n-2}},$$

which is equivalent to

$$\begin{aligned} \delta^2 &> \frac{128\varepsilon^2\psi_0^6(\psi_0^2 - 1)}{\lambda \left(\sqrt{1 + \frac{\lambda}{2\psi_0^2}} - \sqrt{1 - \frac{\lambda}{2\psi_0^2}} \right) (\sqrt{1 + s_0} - \sqrt{1 - s_0})^{n-2}} \\ &\quad + 4\varepsilon\chi \\ &\quad + (\psi_0^4 - 1)(\varepsilon + \chi)^2. \\ \iff \delta^2 &> 8\varepsilon\psi_0^3\chi + 4\varepsilon\chi + (\psi_0^4 - 1)(\varepsilon + \chi)^2. \end{aligned} \tag{13}$$

This bound is again proportional to $(\psi_0^2 - 1)\varepsilon$ and is very similar to the bound given by Inequality (7), made explicit in Inequality (8), but Inequality (13) provides the tougher bound due to the $(\varepsilon + \chi)$ coefficient. \square

We can now provide an upper bound on the length of any canvas edge so that it captures $\text{EV}_{\mathbb{E}}^{-\eta}(p_0)$ and prove Lemma H.2. From Theorem 6.1, we have that if the canvas edge length is bounded as: $e_C < \min\{\mu_0/16, \delta_0^2/64\varepsilon_0\}$, then $V_{\mathbb{E}}(p_0)$ is captured as \mathcal{P} is a δ_0 -power protected (ε_0, μ_0) -net with respect to the Euclidean metric field. As we want to capture $\text{EV}_{\mathbb{E}}(p_0)$, we cannot directly use this result. We have nevertheless obtained the separation between the Voronoi vertices of the eroded Voronoi cell (Equation 11). It is then straightforward to modify the result of Theorem 6.6 by using the separation bound provided in Lemma H.3 instead of the one provided by Lemma C.4. We thus choose

$$e_{C,p_0}^0 = \min \left\{ \frac{\mu}{3}, \frac{\ell_0}{2} \right\}.$$

Remark H.4. *We here ignore the consequences of Appendix G as it only complicates formulas without changing the logical steps.*

We should not forget that we have assumed that g_0 is the Euclidean metric field, which is generally not the case, we must in fact proceed like for the case of a uniform metric field (Theorem H.1):

$$e_{C,p_0} < \min_j \left(\sqrt{\lambda_i} \right) (e_{C,p_0}^0),$$

with $\{\lambda_i\}$ the eigenvalues of g_0 .

Therefore, if the site set satisfies the previous conditions on ε and δ and the canvas is enough for all of its edges to have a length smaller than $e_{C,i}$, then $\text{EV}_{\mathbb{E}}^{-\eta}(p_0)$ is captured, and thus $V_g(p_0)$ is captured, which proves Lemma H.2.

Taking the minimum of all the bounds e_C over all $p_i \in \mathcal{P}$, we obtain an upper bound on the length of the longest canvas edge,

$$e_C = \min_{p_i \in \mathcal{P}} e_{C, p_i},$$

such that all the Voronoi cells are captured.

In all the Lemmas necessary to obtain the local results, we have imposed conditions on ε and δ . Similarly to the bound e_C , the domain-wide bounds on ε and δ are computed from the local values of ε and δ .

Finally, this proves that $\text{Del}_g^d(\mathcal{P}) = \text{Del}_g(\mathcal{P})$ in the general setting (Theorem 6.1), when geodesics are exactly computed.

H.3 Approximate geodesic distance computations

We have so far assumed that geodesics are computed exactly, which is generally not the case in practice. Nevertheless, once the error in the approximation of the geodesic distances is small enough, the computation of the discrete Riemannian Voronoi diagram with approximate geodesic distances can be equivalently seen as the computation of a second discrete Riemannian Voronoi diagram using exact geodesic distances but for a slightly different metric field.

Denote by \tilde{d}_g the geodesic approximation and d_g the exact geodesic distance with respect to the metric field g . Assume that in a small enough neighborhood U (see Lemma B.1), the distances can be related as

$$\left| d_g(p_0, x) - \tilde{d}_g(p_0, x) \right| \leq \xi d_g(p_0, x),$$

where ξ is a function of x that goes to 0 as the sampling parameter ε goes to 0. We can formulate a lemma similar to Lemma E.5 to bound the distance between the same bisectors between sites for the exact and the approximate diagrams.

Lemma H.5. *Let $\psi_0 \geq 1$ be a bound on the metric distortion and g and g' be two Riemannian metric fields on U . Let U be an open neighborhood of $\Omega = \mathbb{R}^2$ that is included in a ball $B_g(p_0, r_0)$, with $p_0 \in U$ and $r_0 \in \mathbb{R}^+$ such that $\forall p \in B(p_0, r_0), \psi(g(p), g'(p)) \leq \psi_0$. Let \mathcal{P}_U be a point set in U and let $p_0 \in \mathcal{P}_U$. Let $\mathcal{P}_U = \{p_i\}$ be a point set in U . Let $\tilde{V}_{p_0, g}$ denote a Voronoi cell with respect to the approximate geodesic distance.*

Suppose that the Voronoi cell $V_g^{+2\xi\tilde{\rho}}(p_0)$ lies in a ball of radius $\tilde{\rho}$ with respect to the metric g , which lies completely in U . Then $\tilde{V}_{p_0, g}$ lies in $V_g^{+2\xi\tilde{\rho}}(p_0)$ and contains $V_g^{-2\xi\tilde{\rho}}(p_0)$.

Proof. Let $\widetilde{\text{BS}}_g(p_0, p_i)$ be the bisector between p_0 and p_i with respect to the approximate geodesic distance. Let $y \in \widetilde{\text{BS}}_g(p_0, p_i) \cap B_g(p_0, \rho)$, where $B_g(p_0, \rho)$ denotes the ball centered at p_0 of radius ρ with respect to the exact geodesic distance. Now $\tilde{d}_g(y, p_0) = \tilde{d}_g(y, p_i)$, and thus

$$\begin{aligned} |d_g(y, p_0)^2 - d_g(y, p_i)^2| &= \left| d_g(y, p_0)^2 - \tilde{d}_g(y, p_0)^2 + \tilde{d}_g(y, p_i)^2 - d_g(y, p_i)^2 \right| \\ &\leq \left| d_g(y, p_0)^2 - \tilde{d}_g(y, p_0)^2 \right| + \left| d_g(y, p_i)^2 - \tilde{d}_g(y, p_i)^2 \right| \\ &\leq 2\xi(d_g(y, p_0)^2 + d_g(y, p_i)^2) \\ &\leq 2\xi\tilde{\rho}^2. \end{aligned}$$

Thus $d_g(y, p_0)^2 \leq d_g(y, p_i)^2 + \tilde{\omega}$ and $d_g(y, p_0)^2 \geq d_g(y, p_i)^2 - \tilde{\omega}$ with $\tilde{\omega} = 2\xi\tilde{\rho}^2$, which gives us the expected result. \square

Denote by $\widetilde{V}(p_0)$ the Voronoi cell of p_0 with the approximate metric. We can then incorporate Lemma H.5 to obtain a result similar to Lemma E.12.

Lemma H.6. *Let U , ψ_0 , g , g_0 and \mathcal{P}_U be defined as in Theorem 5.1. Then we can find η'_0 and ω'_0 such that*

$$EV_{g_0}^{-\eta'_0}(p_0) \subseteq V_{g_0}^{-\omega'_0}(p_0) \subseteq V_g^{-\tilde{\omega}}(p_0) \subseteq \widetilde{V}_g(p_0) \subseteq V_g^{+\tilde{\omega}}(p_0) \subseteq V_{g_0}^{+\omega'_0}(p_0) \subseteq DV_{g_0}^{+\eta'_0}(p_0).$$

These inclusions are illustrated in Figure 15.




Figure 15: Illustration of the different encompassing cells around p_0 in the context of an approximate geodesic distance. The RVDs with respect to g and g_0 are respectively traced in red and black. The dilated Voronoi cells are traced in blue and green. The Voronoi cell $V_{g_0}(p_0)$ is colored in yellow. The cells $DV_{g_0}^{+\eta'_0}$ and $EV_{g_0}^{-\eta'_0}$ are colored in green, and the cells $V_{g_0}^{\pm\omega'_0}(p_0)$ are colored in purple.

The subsequent lemmas and proofs are similar to what was done in the case of exact geodesic computations and we do not explicit them.



**RESEARCH CENTRE
SOPHIA ANTIPOLIS – MÉDITERRANÉE**

2004 route des Lucioles - BP 93
06902 Sophia Antipolis Cedex

Publisher
Inria
Domaine de Voluceau - Rocquencourt
BP 105 - 78153 Le Chesnay Cedex
inria.fr

ISSN 0249-6399